

Drug and Disease Interpretation Learning with Biomedical Entity Representation Transformer

Zulfat Miftahudinov^[0000-0002-8467-4824], Artur Kadurin^[0000-0003-1482-9365],
Roman Kudrin^[0000-0002-9741-0043], and Elena Tutubalina^[0000-0001-7936-0284]

Insilico Medicine Hong Kong, Pak Shek Kok, Hong Kong
zulfat,artur,kudrin,elena@insilico.com
<https://insilico.com/>

Abstract. Concept normalization in free-form texts is a crucial step in every text-mining pipeline. Neural architectures based on Bidirectional Encoder Representations from Transformers (BERT) have achieved state-of-the-art results in the biomedical domain. In the context of drug discovery and development, clinical trials are necessary to establish the efficacy and safety of drugs. We investigate the effectiveness of transferring concept normalization from the general biomedical domain to the clinical trials domain in a zero-shot setting with an absence of labeled data. We propose a simple and effective two-stage neural approach based on fine-tuned BERT architectures. In the first stage, we train a metric learning model that optimizes relative similarity of mentions and concepts via triplet loss. The model is trained on available labeled corpora of scientific abstracts to obtain vector embeddings of concept names and entity mentions from texts. In the second stage, we find the closest concept name representation in an embedding space to a given clinical mention. We evaluated several models, including state-of-the-art architectures, on a dataset of abstracts and a real-world dataset of trial records with interventions and conditions mapped to drug and disease terminologies. Extensive experiments validate the effectiveness of our approach in knowledge transfer from the scientific literature to clinical trials.

Keywords: clinical trials · natural language processing · neural networks · entity linking · medical concept normalization · metric learning · negative sampling · BERT

1 Introduction

The emerging use of neural network architectures in the early-stage of drug discovery has recently resulted in several breakthroughs [50,20]. Later stages of drug development are much more conservative due to the complicated process of clinical trials. The use of state-of-the-art neural network approaches in clinical trials could dramatically speed up the overall drug development process and increase its success rate, thus saving lives.

Clinical trial registers (e.g., ClinicalTrials.gov) contain vast amounts of structured information on how standardized interventions work in a clinical setting.

Despite the existing structure, these registers remain very difficult to harmonize with drug and disease databases using current techniques. This very often results in substantial information losses. The primary cause for this inaccurate harmonization is that in a clinical trial record diseases and interventions are not described with a centralized standardized taxonomy but with a free text. The automatic natural language processing (NLP) methods are promising approaches for the semantic annotation of large volumes of clinical records and for the integration and standardization of biomedical entity mentions to formal concepts. In biomedical research and healthcare, the entity linking problem is known as medical concept normalization (MCN). A source as a knowledge base (KB) contains further information about the concept, such as its preferred name and synonyms, pharmacological profile, and its relationships with other concepts.

Neural architectures have been widely used in recent state-of-the-art models for MCN from user reviews and social media texts [22,49,51,25,31,44]. These studies mostly share limitations regarding a supervised classification framework: binary or multiclass classifiers are trained on a dataset with a narrow subsample of concepts from a specific terminology. In particular, recent models [22,49,51] learn a scoring function measuring the similarity between an entity mention and a concept. The difficulty with these methods is that it is not possible to extract representations describing mentions and concepts separately. In this setup, to retrieve concepts from a particular terminology for a given entity mention, we have to compute all the similarities through the ranking function and sort these scores in descending order. This is impractical if we need to process large corpora of free-form clinical trials, scientific literature, patents in days.

Inspired by metric learning [18,38,16], its usage for multimodal and sentence representation learning [28,37], negative sampling [32], and Bidirectional Encoder Representations from Transformers (BERT) [10], we present a BERT-based neural model for medical concept normalization that directly optimizes the BioBERT representations [23] of entity mentions and concept names itself, rather than classification or ranking layer. We use triplets of free-form entity mention, positive concept names, and randomly sampled concept names as negative examples to train our model. In this work, we consider the zero-shot scenario because it is often the case in the biomedical domain, where there are dozens of concept categories and terminologies. We trained models on annotated pairs of disease or chemical mentions with the corresponding concepts and evaluated on a novel dataset of condition and intervention concepts from clinical trials.

The contributions of this paper can be summarized as follows:

1. We develop a simple and effective model that uses metric learning and negative sampling to obtain entity and concept embeddings. These embeddings were utilized for knowledge transfer between different terminologies. We explore several strategies to select positive and negative samples.
2. We perform extensive experiments of several BERT-based models on a newly annotated dataset of clinical trials in two setups, where each mention is associated with one or more concepts (in-KB) or zero (out-of-KB).

2 Related Work

Our work most closely relates to research in information extraction and semantic textual similarity by directly linking a set of entity mentions and a large set of medical concept names using triplet structures to derive embeddings of entity mentions and concept names that can be compared using semantic similarity. Entity linking of mentions to entries in a knowledge base (KB) is a well-studied area; see a good survey [40]. Research studies in this area assume that there is one knowledge base, such as Wikipedia or Freebase. The KB contains rich text descriptions (from an entity page, for example), hyperlink statistics, and meta-data. This assumption holds for the general domain, but not for the biomedical domain, where diverse terminologies exist for numerous purposes.

2.1 Medical Concept Normalization

Medical concept normalization is usually formulated as a classification or ranking problem with a wide variety of features – syntactic and morphological parsing, dictionaries of medical concepts and their synonyms, distances between raw entity mentions and formal concept names in terms of TF-IDF or word2vec representations [1,12,45,22,9]. MetaMap is one of the most well-known knowledge-based systems for mapping texts to concepts from Unified Medical Language System (UMLS) [3] developed by the US National Library of Medicine (NLM) [1]. This system is based on a linguistic approach using lexical lookup and variants by associating a score with phrases in a sentence. The NLM provides automatic indexing of clinical trials to Medical Subject Headings (MeSH) [6] via the Medical Text Indexer (MTI) [33] based on MetaMap. MTI achieves an F1 measure around 0.55 on the indexing of PubMed abstracts. The most popular open-source supervised system maintained by the NLM is TaggerOne [22]. TaggerOne utilizes semi-Markov models with features and dictionaries to jointly perform entity extraction and normalization tasks.

The works that are the closest to ours and consider synonyms during entity and concept representation learning is Biomedical Named Encoder (BNE) [35] and BioSyn [41]. Sung et al. proposed a BioBERT-based model named BioSyn that maximizes the probability of all synonym representations in the top 20 candidates [41]. BioSyn uses a combination of two scores, sparse and dense, as a similarity function. Sparse scores are calculated on character-level TF-IDF representations to encode morphological information of given strings. Dense scores are defined by the similarity between CLS tokens of a single vector of input in BioBERT. This model achieves state-of-the-art results in disease and chemical mapping over previous works [22,47,35]. Phan et al. presented an encoding framework with new context, concept, and synonym-based objectives [35]. Synonym-based objective enforces similar representations between synonymous names, while concept-based objective pulls the name’s representations closer to its concept’s centroid. However, ranking on these embeddings shows worse results on three sets than TaggerOne.

Our work differs from the studies discussed above in the following important aspects. First, none of these methods have been applied to free-form descriptions of conditions and interventions from clinical trials. Second, evaluation strategies in the mentioned papers are based on train/test splits provided by datasets' authors. We follow the recent *refined* evaluation strategy from [43] on the creation of test sets without duplicates or exact overlaps between the train and test sets. Finally, our dataset includes entity mentions for both in-KB and out-of-KB linking.

2.2 NLP in Clinical Trials Research

While the majority of biomedical research on information extraction primarily focused on scientific literature [17], much less work had been used NLP methods to conduct curation of clinical trial records' fields to advance downstream tasks [11,5,2,4,15,39]. Gayvert et al. [11] proposed an approach for the prediction of the likelihood of toxicity in clinical trials. They selected 108 clinical trials of any phase that were annotated as having failed for toxicity reasons. Then intervention names of each trial were manually mapped to DrugBank [46] concepts to collect molecular weight, polar surface area, and other compounds' properties. In [2], Atal et al. developed a knowledge-based approach to classify entity mentions to disease categories from a Global Burden of Diseases (GBD) cause list. The proposed method uses MetaMap to extract UMLS concepts from trial fields (health condition, public title, and scientific title), link UMLS concepts with ICD10 codes, and classify ICD10 codes to candidate GBD categories. The developed classifier identified GBD categories for 78% of the trials. Li and Lu [26] identified clinical pharmacogenomics (PGx) information from clinical trial records based on dictionaries from a pharmacogenomics knowledge base PharmGKB. Previous studies on clinical trial records, however, have not analyzed the performance of linking of clinical trials to disease and drug concepts, but rather across eligibility criteria (e.g., patient's demographic, disease category) [4,15,39,2,24].

3 Dataset of Clinical Trials

NLM maintains a clinical trial registry data bank ClinicalTrials.gov¹ that contains over 340,000 trials from 214 countries. This database includes comprehensive scientific and clinical investigations in biomedicine [13]. Each trial record provides information about a trial's title, purpose, description, condition, intervention, eligibility, sponsors, etc. Most information from records is described in natural language. In our study, we use publicly available American Association of Clinical Trials (AACT) Database², v. 20200201.

Since there is no off-the-shelf manually annotated dataset for biomedical concept normalization of clinical trials, we built one by selecting 500 trials using the following criteria:

¹ <https://clinicaltrials.gov/>

² <https://www.ctti-clinicaltrials.org/aact-database>

1. A type of clinical study is an interventional study. Participants of interventional studies receive intervention/treatment so that researchers can evaluate the effects of the interventions on biomedical or health-related outcomes [29].
2. Phase of clinical study is defined by U.S. Food and Drug Administration (FDA). There are five phases: Early Phase 1, Phase 1, Phase 2, Phase 3, and Phase 4.
3. Clinical study is associated with one or more interventions of the following types: Biological, Combination Product, Drug.

As a drug terminology source, we use an internal knowledge base that contains 15,532 concept unique identifiers (CUIs), including small molecule drugs, biologics, nutraceuticals, and experimental drugs. As a condition terminology source, we use MeSH v. 20200101. 500 selected trials contain 1075 and 819 entries in the ‘Intervention’ and ‘Condition’ fields respectively. Two annotators with a background in bioinformatics manually annotated each entry. The calculated inter-annotator agreement (IAA) using Kappa was 92.32% for the entire dataset. The disagreement was resolved through mutual consent.

Statistics of annotated texts are summarized in Table 1. 794 out of 1075 non-unique mentions (73.9%) were mapped to one or more drug concepts. 838 (80%) of lower-cased interventions are unique. 804 out of 819 non-unique mentions (98.2%) were mapped to one or more concepts, while there are 638 (78%) lower-cased unique mentions. Interestingly, MeSH concepts linked to conditions belong to several MeSH categories including Diseases [C], Psychiatry and Psychology [F], and Analytical, Diagnostic and Therapeutic Techniques, and Equipment [E]. We note that NLM provided automatically assigned MeSH terms to trials’ interventions. 716 out of 1075 entries (66.6%) were mapped to MeSH terms. Our analysis revealed that mapping from NLM does not include investigational drugs, which are essential for developing new pharmaceutical drugs. Table 2 contains a sample of annotated texts.

4 Model

In this section, we present a neural model for Drug and disease Interpretation Learning with Biomedical Entity Representation Transformer (DILBERT). We address MCN as a retrieval task by fine-tuning the BERT-based network using metric learning [18,38,16], negative sampling [32], specifically, triplet constraints. This idea was successfully applied to learn multimodal embeddings [48,28] and recent sentence embeddings via a sentence-BERT model [37]. Compared to a pair of independent sentences or images, two concept names can have relationships as synonyms, hypernyms, hyponyms, etc., that we consider during the training phase to facilitate the concept ranking task at the retrieval phase.

Let us first recall two terms: *concept* and *concept name*. Following the UMLS Glossary [34], the concept is the fundamental unit of meaning in terminology. It represents a single meaning in any way, whether formal or casual, verbose or abbreviated. Every concept is assigned a unique identifier (CUI). A concept consists of atoms, which are the smallest units of naming. All of the atoms within

Table 1. Statistics of annotated texts.

Mention	#texts	#texts with CUIs	#unique texts	#unique texts with CUIs
Intervention types				
Drug	850	693	671	585
Biological	118	90	102	79
Other	57	4	27	4
Procedure	19	1	16	1
Radiation	11	0	9	0
Device	11	1	11	1
Combination Product	5	3	5	3
Dietary Supplement	2	2	2	2
Diagnostic Test	1	0	1	0
Behavioral	1	0	1	0
Total				
Intervention	1075	794	838	671
Condition	819	804	638	638

a concept are synonymous. The concept name is a string chosen to represent the concept as a whole. It is linked to atoms. Formally, the medical concept normalization task aims to assign each entity mention m a CUI (or predicts that there is no corresponding concept).

Architecture Following denotations proposed by [19], we encode both entity mention m and candidate concept name c into vectors:

$$y_m = \text{red}(T(m)); y_c = \text{red}(T(c)) \quad (1)$$

where T is the transformer that is allowed to update during fine-tuning. $\text{red}(\cdot)$ is a function that reduces that sequence of vectors into one vector. There are two main ways of reducing the output into one representation via $\text{red}(\cdot)$: choose the first output of T (corresponding to the token CLS) or compute the elementwise average over all output vectors to obtain a fixed-size vector. As a pretrained transformer model, we use BioBERT base v1.1. [23]

Scoring The score of a candidate c_i for an entity mention m is given by a distance metric, e.g. Euclidean distance:

$$s(m, c_i) = \|y_m - y_{c_i}\| \quad (2)$$

A noteworthy aspect of the proposed model is its scope: by design, it aims at the cross-terminology mapping of entity mentions to a given lexicon without additional re-training. This approach allows for fast, real-time inference, as all concept names from a terminology can be cached. This is a necessary requirement for processing biomedical documents of different subdomains such as clinical trials, scientific literature and drug labels.

Table 2. Sample of manually annotated trials’ texts.

NCT / Type	Text	Concept
Intervention (with DrugBank CUIs)		
NCT00559975 / Biological	Adjuvanted influenza vaccine combine with CpG7909	Agatolimod sodium (DB15018)
NCT01575756 / Biological	Haemocomplettan® P or Ri-aSTAPTM	Fibrinogen human (DB09222)
NCT00081484 / Drug	epoetin alfa or beta	Erythropoietin (DB00016)
NCT03375593 / Drug	Ibuprofen 600 mg tab	Ibuprofen (DB01050)
NCT01170442 / Drug	vitamin D3 5000 IU	Calcitriol (DB00136)
NCT02493335 / Drug	Placebo orodispersible tablet twice daily	<i>nil (no concept)</i>
Condition (with MeSH CUIs)		
NCT02009605	Squamous Cell Carcinoma of Lung	Carcinoma, Non-Small-Cell Lung (D002289)
NCT04169763	Stage IIC Vulvar Cancer AJCC v8	Vulvar Neoplasms (D014846)

Optimization The network is trained using a triplet objective function. Given a user-generated entity mention m , a positive concept name c_g and a negative concept name c_n , triplet loss tunes the network such that the distance between m and c_g is smaller than the distance between m and c_n . Mathematically, we minimize the following loss function:

$$\max(s(m, c_g) - s(m, c_n) + \epsilon, 0) \quad (3)$$

where ϵ is margin that ensures that c_g is at least ϵ closer to m than c_n . As a scoring metric, we use Euclidean distance or cosine similarity and we set $\epsilon = 1$ in our experiments.

Positive and Negative Sampling Suppose that a pair of the entity mention with the corresponding CUI is given as well as the vocabulary. For positive examples, vocabulary is restricted to the concepts that have the same CUI as a mention. Multiple positive concept names could be explained by the presence of synonyms in the vocabulary. Negative sampling [32] uses the rest part of the vocabulary. We explore several strategies to select positive and negative samples for a training pair (entity mention, CUI):

1. **random sampling**: we sample several concept names with the same CUI as positive examples and random negatives from the rest of the vocabulary;
2. **random + parents**: we sample k concept names from the concept’s parents in addition to positive and negative names gathered with the random sampling strategy;
3. **re-sampling**: using a model trained with random sampling, we identify positives and *hard* negatives via the following steps: (i) encode all mentions

and concept names found in training pairs using the current model, (ii) select positives with the same CUI, which are closest to a mention, (iii) for each mention, retrieve the most similar k concept names (i.e., its nearest neighbors) and select all names that are ranked above the correct one for the mention as negative examples. We follow this strategy from [14];

4. **re-sampling + siblings**: we modify the re-sampling strategy by using k concept names from the concept’s siblings as negatives.

Inference At inference time, the representation for all concept names can be precomputed and cached. The inference task is then reduced to finding the closest concept name representation to entity mention representation in a common embedding space.

5 Experiments

We evaluate our model DILBERT and compare it to the state-of-the-art methods using (i) a publicly available benchmark BioCreative V CDR Disease & Chemical [27], (ii) our dataset of clinical trials named CT Condition & Intervention. The statistics of the two datasets are summarized in Table 3.

5.1 Datasets

BioCreative V CDR [27] introduces a challenging task for the extraction of chemical-disease relations (CDR) from PubMed abstracts. Disease and chemical mentions are linked to the MEDIC [8] and CTD [7] dictionaries, respectively. We utilize the CTD chemical dictionary (v. November 4, 2019) that consists of 171,203 CUIs and 407,247 synonyms, and the MEDIC lexicon (v. July 6, 2012) that contains 11,915 CUIs and 71,923 synonyms.

According to the BioCreative V CDR annotation guidelines, the annotators used two MeSH branches to annotate entities: (i) “Diseases” [C], including signs and symptoms, (ii) “Drugs and Chemicals” [D]. The terms “drugs” and “chemicals” are often used interchangeably. Annotators annotated chemical nouns convertible to single atoms, ions, isotopes, pure elements and molecules (e.g., calcium, lithium), class names (e.g., steroids, fatty acids), small biochemicals, synthetic polymers.

As shown in [43], the CDR dataset contains a high amount of mention duplicates and overlaps between official sets. In order to obtain more realistic results, we evaluate models on preprocessed official and *refined* CDR test sets from [43].

For the preprocessing of the clinical trial data, we use heuristic rules to split the composite mentions into separate mentions (e.g., *combination of ribociclib + capecitabine* into *ribociclib* and *capecitabine*) by considering each mention containing “combination”, “combine”, “combined”, “plus”, “vs” or “+” as composite. We process all characters to lowercase forms and remove the punctuation for both mentions and synonyms.

Table 3. Statistics of the datasets used in the experiments. Two sets of annotated clinical trials’ fields are marked with ‘CT’.

	CDR Disease	CDR Chem	CT Condition	CT Intervention
domain	abstracts	abstracts	clinical trials	clinical trials
entity type	disease	chemicals	conditions	drugs
terminology	MEDIC	CTD Chemicals	MeSH	in-house dict.
entity level statistics				
% numerals	0.11%	7.32%	7.69%	25.3%
% punctuation	1.21%	0.07%	14.28%	24.83%
avg. len	14.88	11.27	17.92	21.68
number of pre-processed entity mentions				
train set	4,182	5,203	-	-
dev set	4,244	5,347	100	100
test set	4,424	5,385	719	975
number of pre-processed entity mentions after removal of duplicates from test set				
<i>refined</i> test	657 (14.9%)	425 (7.9%)	638 (77.89%)	838 (77.95%)

It is assumed that each entity mention in the CDR corpus has a valid concept in the terminology, which is referred as in-KB evaluation in the entity linking task. In contrast with the CDR sets, 26% and 1.8% of intervention and condition mentions in the CT dataset are not appeared in terminologies, respectively. In Section 5.4, we investigate different strategies for the out-of-KB prediction (i.e. *nil* prediction) on clinical trials’ texts.

5.2 Baseline Methods

We compare our proposed method with the following methods.

BioBERT ranking This is a baseline model that used the BioBERT model for encoding mention and concept representations. Each entity mention or concept name is firstly passed through BioBERT (we use the average over all outputs of BERT) and then through a mean pooling layer to yield a fixed-sized vector. The inference task is then reduced to finding the closest concept name representation to entity mention representation in a common embedding space. We use the Euclidean distance as the distance metric. The nearest concept names are chosen as top-k concepts for entities. We use the publicly available code provided by [43] at <https://github.com/insilicomedicine/Fair-Evaluation>.

BioSyn BioSyn [41] is a recent state-of-the-art model that utilizes the synonym marginalization technique and the iterative candidate retrieval. The model uses two similarity functions based on sparse and dense representations, respectively. The sparse representation encodes the morphological information of given strings via TF-IDF, the dense representation encodes the semantic information gathered from BioBERT. For reproducibility, we use the publicly available code provided by the authors at <https://github.com/dmis-lab/BioSyn>. We follow the default parameters of BioSyn as in [41]: the number of top candidates k is 20, the

mini-batch size is 16, the learning rate is $1e-5$, the dense ratio for the candidate retrieval is 0.5, 20 epochs for training.

5.3 Experimental Setup

We experiment with BioBERT_{base} v1.1 with 12 heads, 12 layers, 768 hidden units per layer, and a total of 110M parameters. Epsilon, the number of positive and negative examples, and distance metric were chosen optimally on dev sets. We choose $red(\cdot)$ to be the average over all outputs of BERT. We have evaluated different epsilons starting from 0.5 up to 4.0 with 0.5 step for Euclidean distance metric, for cosine distance from 0.05 up to 0.3 with 0.05 step. These experiments have quite similar results. We have evaluated a number of positive and negative examples. For positives, we iterated over values from 15 to 35, for negatives from 5 to 15. We found that the optimal is to sample 30 positive examples and 5 negative examples per mention. For the random + parents strategy, we evaluated the number of names of concept’s parents from 1 to 5. Similarly, we evaluated the number of names of concept’s siblings from 1 to 5. We found that hard negative sampling (with siblings) achieves the same optima as random negative sampling. The highest metrics are achieved at 5 concept names of the concept’s parents on the CT Condition and CDR Chemical sets. The highest accuracy is achieved at 2 names of the concept’s parents on other sets. As a result, we trained the DILBERT model with Euclidean distance and the following parameters: batch size is equal to 48, learning rate was set to $1e-5$, epsilon to 1.0.

We evaluate this solution in information retrieval (IR) scenario, where the goal is to find within a dictionary of concept names and their identifiers the top- k concepts for every entity mention in texts. In particular, we use the top- k accuracy as an evaluation metric, following the previous works [42,36,47,35,41,43]. Let $Acc@k$ be 1 if a right CUI is retrieved at rank k , otherwise 0. All models are evaluated with $Acc@1$. For composite entities, we define $Acc@k$ as 1 if each prediction for a single mention is correct.

5.4 Out-of-KB Cases in Clinical Trials

To deal with *nil* predictions in clinical trials, we apply three different strategies for the selection of a threshold value. Namely, the intervention or condition mention is considered out of KB if the nearest candidate has a larger distance than a threshold value. Our first strategy is to set the threshold equal to the minimum distance of false-positive (FP) cases. In this case, we consider a mention mapped to a concept by our model but having no appearance in the terminology. Our second strategy set the threshold to the maximum distance of true-positive (TP) cases. The third strategy uses a weighted average of the first two threshold values. The proportion of TP cases used as a weight for the first strategy’s threshold, the proportion of TP cases used as a weight for the second strategy’s threshold. We tested three strategies on the dev set which containing 100 randomly selected mentions and evaluated the selected threshold values on the test set. This procedure was repeated 20 times. For intervention normalization,

Table 4. Out-of-domain performance of the proposed DILBERT model and baselines in terms of Acc@1 on the *refined* test set of clinical trials (CT).

Model	CT Condition		CT Intervention	
	single concept	full set	single concept	full set
BioBERT ranking	72.60	71.74	78.67	74.57
BioSyn	86.36	-	86.29	-
DILBERT, random sampling	85.73	84.85	90.23	88.37
DILBERT, random + 2 parents	86.74	86.36	90.53	87.94
DILBERT, random + 5 parents	87.12	86.74	89.54	87.15
DILBERT, resampling	85.22	84.63	89.83	87.28
DILBERT, resampling + 5 siblings	84.84	84.26	89.26	86.23

Table 5. In-domain performance of the proposed DILBERT model in terms of Acc@1 on the *refined* test set of the Biocreative V CDR corpus.

Model	CDR Disease	CDR Chemical
BioBERT ranking	66.4	80.7
BioSyn	74.1	83.8
DILBERT, random sampling	75.5	81.4
DILBERT, random + 2 parents	75.0	81.2
DILBERT, random + 5 parents	73.5	81.4
DILBERT, resampling	75.8	83.3
DILBERT, resampling + 5 siblings	75.3	82.1

the first strategy showed an average accuracy of 79.41 with std of 3.5; second – accuracy of 71.77 and std of 3.5; third – accuracy of 85.73, std of 1.3.

5.5 Results and Discussion

We investigate the effectiveness of transferring concept normalization from the general biomedical domain to the clinical trial domain. We trained DILBERT and BioSyn models on the CDR Disease and CDR Chemical train sets, respectively, for linking clinical conditions and interventions.

Table 4 presents the performance of the DILBERT models compared to BioSyn and BioBERT ranking on the datasets of clinical trials. We test the DILBERT model’s transferability on two sets of interventions and conditions where each mention is associated with one concept only (see ‘single concept’ columns). We evaluate the model on test sets with all mentions, including single concepts, composite mentions, and out-of-KB cases (see ‘full set’ columns). In Table 5, we present in-domain results of models evaluated on the CDR data. In all our experiments when comparing DILBERT and BioSyn models, we use paired McNemar’s test [30] with a confidence level at 0.05 to measure statistical significance.

Several observations can be made based on Tables 4 and 5. First, DILBERT outperformed BioSyn and BioBERT ranking on three sets staying on par with

BioSyn on the CDR Chemical test set. Adding randomly sampled positive examples from parent-child relationships gives a statistically significant improvement in 1-2% on the CT Condition set while staying on par with random sampling on interventions. To our surprise, hard negative mining produces performance gains on one of four sets only, which includes chemicals. Second, we compare results on refined test sets with results on the CDR corpus’s official test set. We observe the significant decrease of Acc@1 from 93.6% to 75.8% and from 95.8% to 83.8% for DILBERT on disease and chemical mentions, respectively. Third, DILBERT models obtained higher results on test sets with single concepts. Models achieve much higher performance for the normalization of interventions rather than conditions. The DILBERT model achieves a statistically significant improvement compared to the BioSyn model on the interventions dataset. The error analysis on the CDR Disease set showed that models with random negative sampling incorrectly maps 39 out of 147 mentions to the correct concept’s parent. We observe that some mentions are mapped to the gold concept’ child for the models trained by re-sampling+siblings sampling.

Inference Time Efficiency and Deployment Our model uses the FAISS library [21] with GPU support for fast nearest neighbor search by comparing vectors with Euclidean distance. Embeddings of all terminologies’ concepts are indexed. We profiled retrieval speed on a server with Intel Xeon CPU E5-2660 2.00GHz and 256GB memory. First, we precomputed all embeddings for all concepts (500 thousand). On a single Nvidia TITAN X GPU, it takes about 7 minutes to compute all embeddings. Given that all embeddings are indexed on Nvidia TITAN X GPU using IndexFlatL2 index type. To obtain top candidates for 10 million queries, it requires approximately 3 hours.

6 Conclusion

We studied the task of drug and disease normalization for clinical trials, using a newly created dataset of 500 interventional studies with 1075 intervention mentions and 819 condition mentions. We designed a triplet-based metric learning model named DILBERT that optimizes to pull pairs of mention and concept BioBERT representations closer than negative samples. We investigated strategies to obtain random and hard positive and negative examples using parent-child (i.e., broader-narrower) relationships between biomedical concepts. We performed experiments on in-KB and out-of-KB (*nil*) linking of mentions from the scientific domain to the clinical domain in a zero-shot setting. DILBERT shows better transfer capabilities for disease- and drug-related mentions compared to other state-of-the-art models. In future work, we plan to investigate taxonomy induction evaluation metrics and the normalization of protein/gene mentions.

Acknowledgements Research on academic corpora was carried out by Z.M. and supported by RFBR, project no. 19-37-90074.

References

1. Aronson, A.R.: Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. In: Proceedings of the AMIA Symposium. p. 17. American Medical Informatics Association (2001)
2. Atal, I., Zeitoun, J.D., Névéol, A., Ravaud, P., Porcher, R., Trinquart, L.: Automatic classification of registered clinical trials towards the global burden of diseases taxonomy of diseases and injuries. *BMC bioinformatics* **17**(1), 392 (2016)
3. Bodenreider, O.: The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research* **32**(suppl.1), D267–D270 (2004)
4. Boland, M.R., Miotto, R., Gao, J., Weng, C.: Feasibility of feature-based indexing, clustering, and search of clinical trials. *Methods of information in medicine* **52**(05), 382–394 (2013)
5. Brown, A.S., Patel, C.J.: A standard database for drug repositioning. *Scientific data* **4**(1), 1–7 (2017)
6. Coletti, M.H., Bleich, H.L.: Medical subject headings used to search the biomedical literature. *Journal of the American Medical Informatics Association* **8**(4), 317–323 (2001)
7. Davis, A.P., Grondin, C.J., Johnson, R.J., Sciaky, D., McMorran, R., Wieggers, J., Wieggers, T.C., Mattingly, C.J.: The comparative toxicogenomics database: update 2019. *Nucleic acids research* **47**(D1), D948–D954 (2019)
8. Davis, A.P., Wieggers, T.C., Rosenstein, M.C., Mattingly, C.J.: Medic: a practical disease vocabulary used at the comparative toxicogenomics database. *Database* **2012** (2012)
9. Dermouche, M., Looten, V., Flicoteaux, R., Chevret, S., Velcin, J., Taright, N.: ECSTRA-INSERM@ CLEF eHealth2016-task 2: ICD10 code extraction from death certificates. *CLEF* (2016)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 4171–4186 (2019)
11. Gayvert, K.M., Madhukar, N.S., Elemento, O.: A data-driven approach to predicting successes and failures of clinical trials. *Cell chemical biology* **23**(10), 1294–1301 (2016)
12. Ghiasvand, O., Kate, R.J.: Uwm: Disorder mention extraction from clinical text using crfs and normalization using learned edit distance patterns. In: *SemEval@ COLING*. pp. 828–832 (2014)
13. Gill, S.K., Christopher, A.F., Gupta, V., Bansal, P.: Emerging role of bioinformatics tools and software in evolution of clinical research. *Perspectives in clinical research* **7**(3), 115 (2016)
14. Gillick, D., Kulkarni, S., Lansing, L., Presta, A., Baldrige, J., Ie, E., Garcia-Olano, D.: Learning dense representations for entity retrieval. In: Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL). pp. 528–537 (2019)
15. Hao, T., Rusanov, A., Boland, M.R., Weng, C.: Clustering clinical trials with similar eligibility criteria features. *Journal of biomedical informatics* **52**, 112–120 (2014)
16. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: International Workshop on Similarity-Based Pattern Recognition. pp. 84–92. Springer (2015)

17. Huang, C.C., Lu, Z.: Community challenges in biomedical text mining over 10 years: success, failure and the future. *Briefings in bioinformatics* **17**(1), 132–144 (2015)
18. Huang, P.S., He, X., Gao, J., Deng, L., Acero, A., Heck, L.: Learning deep structured semantic models for web search using clickthrough data. In: *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. pp. 2333–2338 (2013)
19. Humeau, S., Shuster, K., Lachaux, M.A., Weston, J.: Poly-encoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring. *CoRR* abs/1905.01969. External Links: Link Cited by **2**, 2–2 (2019)
20. Ivanenkov, Y., Zhavoronkov, A., Yamidanov, R., Osterman, I., Sergiev, P., Aladinskiy, V., Aladinskaya, A., Terentiev, V., Veselov, M., Ayginin, A., et al.: Identification of novel antibacterials using machine-learning techniques. *Frontiers in pharmacology* **10**, 913 (2019)
21. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734* (2017)
22. Leaman, R., Lu, Z.: Taggerone: joint named entity recognition and normalization with semi-markov models. *Bioinformatics* **32**(18), 2839–2846 (2016)
23. Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J.: Biobert: pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* (09 2019)
24. Leveling, J.: Patient selection for clinical trials based on concept-based retrieval and result filtering and ranking. In: *TREC* (2017)
25. Li, H., Chen, Q., Tang, B., Wang, X., Xu, H., Wang, B., Huang, D.: Cnn-based ranking for biomedical entity normalization. *BMC bioinformatics* **18**(11), 79–86 (2017)
26. Li, J., Lu, Z.: Systematic identification of pharmacogenomics information from clinical trials. *Journal of biomedical informatics* **45**(5), 870–878 (2012)
27. Li, J., Sun, Y., Johnson, R.J., Sciaky, D., Wei, C.H., Leaman, R., Davis, A.P., Mattingly, C.J., Wiegers, T.C., Lu, Z.: Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database* **2016** (2016)
28. Liu, Y., Guo, Y., Bakker, E.M., Lew, M.S.: Learning a recurrent residual fusion network for multimodal matching. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4107–4116 (2017)
29. Lo, B.: *Sharing clinical trial data: maximizing benefits, minimizing risk*, vol. 313. American Medical Association (2015)
30. McNemar, Q.: Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* **12**(2), 153–157 (1947)
31. Miftahutdinov, Z., Tutubalina, E.: Deep neural models for medical concept normalization in user-generated texts. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. pp. 393–399 (2019)
32. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*. pp. 3111–3119 (2013)
33. Mork, J.G., Jimeno-Yepes, A., Aronson, A.R.: The.nlm medical text indexer system for indexing biomedical literature. In: *BioASQ@ CLEF* (2013)
34. *NLM: Umls glossary* (2016), http://www.nlm.nih.gov/research/umls/new_users/glossary.html

35. Phan, M.C., Sun, A., Tay, Y.: Robust representation learning of biomedical names. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 3275–3285 (2019)
36. Pradhan, S., Elhadad, N., Chapman, W.W., Manandhar, S., Savova, G.: Semeval-2014 task 7: Analysis of clinical text. In: SemEval@ COLING. pp. 54–62 (2014)
37. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). pp. 3973–3983 (2019)
38. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
39. Sen, A., Goldstein, A., Chakrabarti, S., Shang, N., Kang, T., Yaman, A., Ryan, P.B., Weng, C.: The representativeness of eligible patients in type 2 diabetes trials: a case study using gist 2.0. *Journal of the American Medical Informatics Association* **25**(3), 239–247 (2018)
40. Shen, W., Wang, J., Han, J.: Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* **27**(2), 443–460 (2014)
41. Sung, M., Jeon, H., Lee, J., Kang, J.: Biomedical entity representations with synonym marginalization. *arXiv preprint arXiv:2005.00239* (2020)
42. Suominen, H., Salanterä, S., Velupillai, S., Chapman, W.W., Savova, G., Elhadad, N., Pradhan, S., South, B.R., Mowery, D.L., Jones, G.J., et al.: Overview of the share/clef ehealth evaluation lab 2013. In: International Conference of the Cross-Language Evaluation Forum for European Languages. pp. 212–231. Springer (2013)
43. Tutubalina, E., Kadurin, A., Miftahutdinov, Z.: Fair evaluation in concept normalization: a large-scale comparative analysis for bert-based models. In: Proceedings of the 28th International Conference on Computational Linguistics. pp. 6710–6716 (2020)
44. Tutubalina, E., Miftahutdinov, Z., Nikolenko, S., Malykh, V.: Medical concept normalization in social media posts with recurrent neural networks. *Journal of biomedical informatics* **84**, 93–102 (2018)
45. Van Mulligen, E., Afzal, Z., Akhondi, S.A., Vo, D., Kors, J.A.: Erasmus MC at CLEF eHealth 2016: Concept recognition and coding in French texts. CLEF (2016)
46. Wishart, D.S., Knox, C., Guo, A.C., Shrivastava, S., Hassanali, M., Stothard, P., Chang, Z., Woolsey, J.: Drugbank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic acids research* **34**(suppl_1), D668–D672 (2006)
47. Wright, D., Katsis, Y., Mehta, R., Hsu, C.N.: Normco: Deep disease normalization for biomedical knowledge base construction. In: Automated Knowledge Base Construction (2019), <https://openreview.net/forum?id=BJerQWcp6Q>
48. Wu, P., Hoi, S.C., Xia, H., Zhao, P., Wang, D., Miao, C.: Online multimodal deep similarity learning with application to image retrieval. In: Proceedings of the 21st ACM international conference on Multimedia. pp. 153–162 (2013)
49. Zhao, S., Liu, T., Zhao, S., Wang, F.: A neural multi-task learning framework to jointly model medical named entity recognition and normalization. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 817–824 (2019)
50. Zhavoronkov, A., Ivanenkov, Y.A., Aliper, A., Veselov, M.S., Aladinskiy, V.A., Aladinskaya, A.V., Terentiev, V.A., Polykovskiy, D.A., Kuznetsov, M.D., Asadulaev, A., et al.: Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature biotechnology* **37**(9), 1038–1040 (2019)

51. Zhu, M., Celikkaya, B., Bhatia, P., Reddy, C.K.: Latte: Latent type modeling for biomedical entity linking. arXiv preprint arXiv:1911.09787 (2019)