

Zero-shot Medical Entity Retrieval without Annotation: Learning From Rich Knowledge Graph Semantics

Luyang Kong

Amazon AI

luyankon@amazon.com

Christopher Winestock

Amazon AI

winestoc@amazon.com

Parminder Bhatia

Amazon AI

parmib@amazon.com

Abstract

Medical entity retrieval is an integral component for understanding and communicating information across various health systems. Current approaches tend to work well on specific medical domains but generalize poorly to unseen sub-specialties. This is of increasing concern under a public health crisis as new medical conditions and drug treatments come to light frequently. Zero-shot retrieval is challenging due to the high degree of ambiguity and variability in medical corpora, making it difficult to build an accurate similarity measure between mentions and concepts. Medical knowledge graphs (KG), however, contain rich semantics including large numbers of synonyms as well as its curated graphical structures. To take advantage of this valuable information, we propose a suite of learning tasks designed for training efficient zero-shot entity retrieval models. Without requiring any human annotation, our knowledge graph enriched architecture significantly outperforms common zero-shot benchmarks including BM25 and Clinical BERT with 7% to 30% higher recall across multiple major medical ontologies, such as UMLS, SNOMED and ICD-10.

1 Introduction

Entity retrieval is the task of linking mentions of named entities to concepts in a curated knowledge graph (KG). This allows medical researchers and clinicians to search medical literature easily using standardized codes and terms to improve patient care. Training an effective entity retrieval system often requires high quality annotations, which are expensive and slow to produce in the medical domain. It is therefore not feasible to annotate enough data to cover the millions of concepts in a medical KG, and difficult to adapt quickly enough to those newly appeared medical conditions and drug treatments under a public health crisis. Hence, a robust

medical entity retrieval system is expected to have decent performance in a zero-shot scenario.

Zero-shot retrieval is challenging due to the complexity of medical corpora - large numbers of ambiguous terms, copious acronyms and synonymous terms. It is difficult to build an accurate similarity measure which can detect the true relatedness between a mention and a concept even when their surface forms differ greatly.

Early entity retrieval systems use string matching methods such as exact match, approximate match (Hanisch et al., 2005) and weighted keyword match e.g. BM25 (Wang et al., 2020). Although annotated training data is not required, such systems typically lack the ability to handle synonyms and paraphrases with large surface form differences. In recent years, large scale pretraining (Devlin et al., 2019) has been widely adopted in the medical domain such as Clinical BERT (Alsentzer et al., 2019) and BioBERT (Lee et al., 2019). Agarwal et al. (2019) also integrates graph structure information during pretraining. Most of them, however, require a finetuning step on annotated training data (Wu et al., 2020) before being applied to entity retrieval.

As an alternative to manually annotating a corpus, the rich semantics inside a KG itself can be utilized (Chang et al., 2020). One important entry is the **synonym**, whereby two medical terms may be used interchangeably. In addition, the **graphical structure** of a KG contains information on how concepts are related to each other and so can be used as another valuable resource for building an effective similarity measure. We therefore design *synonym-based tasks* and *graph-based tasks* to mine a medical KG. Trained with our proposed tasks, a simple Siamese architecture significantly outperforms common zero-shot benchmarks across multiple major medical ontologies including UMLS, SNOMED and ICD10.

Our contributions are as follows. (1) We pro-

pose a framework which allows the information in medical KGs to be incorporated into entity retrieval models, thereby enabling robust zero-shot performance without the need of human annotations. (2) We apply the framework to major medical ontologies and conduct extensive experiments to establish the effectiveness of our framework. (3) When annotations are available, we show that the proposed framework can be easily plugged into an existing supervised approach and in so doing, deliver consistent improvements.

2 Formulation

Entity retrieval. Given a *mention* m and a *concept* $c \in KG = \{c_1, c_2, \dots, c_n\}$, the goal is to learn a similarity measurement $S(m, c)$, so that the most relevant concept is assigned the highest score. A concept is also referred to as a *node* in a KG. We use them interchangeably below.

Zero-shot entity retrieval. We examine two zero-shot scenarios: 1) *zero-shot on mentions only*, which assumes unseen mentions but allows seen concepts at test time. 2) *zero-shot on mentions and concepts*, which assumes both to be unseen at test time.

3 Model Architecture

Siamese architecture. Mention m and concept c are firstly embedded into vectors, using a shared function T : $e_m = T(m)$, $e_c = T(c)$. T is also referred to as an encoder, for which we use the Transformer (Vaswani et al., 2017) encoder in this work. Similarity between a mention and a concept is then measured as the inner product: $S(m, c) = \langle e_m, e_c \rangle$.

Optimization. Assume model parameter is θ . We use in-batch negatives for optimization. Loss function for a batch of size B is defined as mean negative log likelihood:

$$L = -\frac{1}{B} \sum_{i=1}^B \log(P((m_i, c_i)|\theta))$$

where the conditional probability of each mention-concept pair (m_i, c_i) in the batch is modeled as a softmax:

$$P((m_i, c_i)|\theta) = \frac{\exp(S_\theta(m_i, c_i))}{\sum_{j=1}^B \exp(S_\theta(m_j, c_j))}$$




Figure 1: ICD-10 synonym-based task defined at an example node

4 Learning Task

We design our learning tasks by constructing mention-concept pairs (m, c) . The goal is to capture multiple layers of semantics from a KG by leveraging its unique structure. Since each structure implies its own measure of similarity, we design learning tasks by finding very similar or closely related textual descriptions and use them to construct (m, c) pairs. We define two major types of tasks: *synonym-based* tasks and *graph-based* tasks. These are illustrated below for three major medical KG: ICD-10, SNOMED and UMLS.

4.1 ICD-10

The 10th version of the International Statistical Classification of Diseases, Clinical Modification (ICD-10) is one of the most widely used terminology systems for medical conditions. It contains over 69K concepts, organized in a tree structure of parent-child relationships.

Synonym-based task. In ICD-10, a child node is a more specific medical condition compared to its parent (e.g. R07.9 Chest pain, unspecified is a child of R52 Pain, unspecified). Each node N_i has three sections: The *Title* section contains a specifier (e.g. Chest) of the title of the parent (e.g. Pain), therefore their concatenation gives the full concept description (e.g. Chest Pain). We denote it by $N_i^{TitleConcatenation}$. The *Code* section contains an ICD-10 code and its formal medical definition, denoted by $N_i^{CodeDescription}$. The *SeeAlso* section contains a similar concept, denoted by $N_i^{SeeAlso}$.

These three sections describe the same medical

condition with different surface forms, therefore we define the ICD-10 synonym-based task as:

$$\mathbf{m} = N_i^L, \mathbf{c} = N_i^R$$

$$N_i^L, N_i^R \in \{N_i^{TitleConcatenation}, N_i^{CodeDescription}, N_i^{SeeAlso}\}, L \neq R$$

We illustrate it with an example in Figure 1.

Graph-based task. To incorporate the semantics of parent-child relationships into learning, we define ICD-10 graph-based task as:

$$\mathbf{m} = N_i^{CodeDescription}, \mathbf{c} = N_j^{CodeDescription}$$

$$N_i.is_parent(N_j)$$

4.2 SNOMED

Systematized Nomenclature of Medicine – Clinical Terms (SNOMED) is a standardized clinical terminology used for the electronic exchange of clinical health information with over 360K active concepts.

Synonym-based task. Each node N_i in SNOMED has multiple synonymous descriptions $\{l_i^1, l_i^2, \dots, l_i^d\}$, with l_i^1 as the main description. We therefore define SNOMED synonym-based task as:

$$\mathbf{m} = l_i^p, \mathbf{c} = l_i^q, p > q$$

$\frac{d*(d-1)}{2}$ unique (\mathbf{m}, \mathbf{c}) pairs are constructed at each node.

Graph-based task. SNOMED is a directed graph with 107 possible relationship types (e.g. `is_a`, `finding_site`, `relative_to`). A direct connection between two nodes is likely to imply a certain degree of similarity, thus we define the SNOMED graph-based task as:

$$\mathbf{m} = l_i^1, \mathbf{c} = l_j^1$$

$$N_i.is_connected(N_j)$$

4.3 UMLS

The Unified Medical Language System (UMLS) is a compendium of a large number of curated biomedical vocabularies with over 1MM concepts. UMLS has almost the same structure as SNOMED, therefore we define the **synonym-based task** and **graph-based task** in a similar fashion to that of SNOMED.

For each task mentioned above, the (\mathbf{m}, \mathbf{c}) pairs generated at each node are combined and split into train and dev in a 80:20 ratio. We also define a **comb task**, where all the tasks are firstly down-sampled to equal sizes and then combined. A summary can be found in Table 1.

KG	Task Type	Train	Dev
ICD-10	syn	113K	28K
	graph	33K	8K
SNOMED	syn	1.4M	374K
	graph	955K	238K
UMLS	syn	27M	7M
	graph	7M	2M
Comb (by down-sampling)		198K	488K

Table 1: **Task Description:** Number of (\mathbf{m}, \mathbf{c}) pairs in train and dev for all tasks.

Dataset	Split	KG	Test size
MedM.	-	UMLS	66,572
	SG		4,350
COMETA	SS	SNOMED	4,369
	ZG		3,995
	ZS		4,283
3DNotes	ICD	ICD-10	5,742
	SN	SNOMED	7,521

Table 2: **Test Set Size.**

5 Datasets

We include three datasets in zero-shot evaluations.

MedMention (Mohan and Li, 2019) is a publicly available corpus of 4,392 PubMed¹ abstracts with biomedical entities annotated with UMLS concepts. **COMETA** (Basaldella et al., 2020) is one of the largest public corpora of social media data with SNOMED annotations. It provides four train, dev, test splits: Stratified-General (SG), Stratified-Specific (SS), Zeroshot-General (ZG), Zeroshot-Specific (ZS). We also use a de-identified corpus of dictated doctor’s notes named **3DNotes** (Zhu et al., 2020). It has two sets of annotations: one with ICD-10 (ICD split), another with SNOMED (SN split). The annotation follows the i2b2 challenge (Uzuner et al., 2011) guidelines.

Zero-shot performance is evaluated on the corresponding test sets. We report sizes of the test sets in Table 2.

6 Experiments

6.1 Experimental Settings

We train a model using each task and evaluate them across all test sets. Performance is compared against two common zero-shot entity retrieval benchmarks: BM25, and BERT encoder followed by inner-product similarity. For the latter, we tested multiple pre-trained versions including BERT_{base}, Clinical BERT and BioBERT.

¹<https://www.ncbi.nlm.nih.gov/pmc/>

Dataset	Split	KG	Siamese architecture trained with KG learning tasks (ours)								
			BM25	Clinical BERT	ICD-10		SNOMED		UMLS		Comb
					Syn	Graph	Syn	Graph	Syn	Graph	
MedM.	-	UMLS	.04(.17)	.10(.30)	.31(.58)	.31(.56)	.32(.55)	.33(.61)	.32(.53)	.30(.57)	.32(.60)
COMETA	SG	SNOMED	.02(.10)	.01(.06)	.30(.52)	.30(.48)	.43(.65)	.37(.58)	.33(.50)	.32(.54)	.37(.58)
	SS		.02(.11)	.01(.06)	.28(.51)	.28(.47)	.41(.62)	.36(.56)	.31(.48)	.31(.52)	.35(.56)
	ZG		.02(.12)	.01(.07)	.32(.57)	.32(.54)	.47(.71)	.39(.61)	.36(.55)	.33(.57)	.40(.62)
	ZS		.02(.10)	.01(.07)	.30(.52)	.29(.47)	.40(.64)	.35(.57)	.31(.49)	.29(.53)	.35(.57)
3DNotes	ICD	ICD-10	.05(.22)	.11(.17)	.28(.54)	.23(.46)	.20(.45)	.20(.52)	.18(.39)	.21(.53)	.30(.54)
	SN	SNOMED	.07(.20)	.01(.05)	.20(.50)	.18(.45)	.38(.63)	.25(.61)	.25(.49)	.29(.55)	.34(.59)

Table 3: Retrieval performance R@1(25). Siamese architecture trained with our tasks are shown to significantly outperform benchmarks. Evaluation for *zero-shot on mentions only* is highlighted in **bold** the rest belongs to *zero-shot on mentions and concepts*. The former enjoys a bigger gain as expected.

Mention	Gold Concept	Syn	Graph
shortness of breath	dyspnea (finding)	✓	✗
GI hemorrhage	gastrointestinal hemorrhage (disorder)	✓	✗
coronary structure	coronary artery (body structure)	✗	✓
heart	heart structure (body structure)	✗	✓

Table 4: Prediction error of the model trained with SNOMED tasks evaluated on 3DNotes-SN.

Hyperparameters. For our Siamese architecture, the transformer encoder is initialized with BERT_{base}. We use the BertAdam optimizer with a batch size of 128, the initial learning rate of 3×10^{-5} , warm-up ratio of 0.02, max epochs of 50, followed by a linear learning rate decay.

Evaluation metrics. Top- k retrieval recalls (R@1, R@25) are used as metrics. We also assume that each mention has a valid gold concept in the KG.

6.2 Results

We report overall results in Table 3. Clinical BERT consistently outperforms the other pre-trained counterparts, which are therefore omitted. For evaluations of *zero-shot on mentions only* (e.g. UMLS tasks evaluated on MedMention which is UMLS annotated), we observe 12% to 45% gain for R@1 compared to benchmarks. For evaluations of *zero-shot on mentions and concepts* (e.g. UMLS tasks evaluated on COMETA which is SNOMED annotated), 7% to 30% higher R@1 is observed. *Comb* task has the most balanced performance gains across all datasets.

6.3 Analysis and Discussion

Task comparison. To further understand the difference between synonym-based tasks and graph-based tasks, we illustrate qualitative examples in Table 4. A model trained using the synonym task




Figure 2: R@25. Using the *comb* task as an auxiliary to the primary supervised loss, the model gains 1% in-domain (3DNotes-SN) improvement, 8% zero-shot (COMETA-ZS) improvement.

makes better predictions for scenarios involving medical synonyms and acronym (lines 1, 2). A model trained using the graph task performs better when mention and concept have an *is_a* relationship (lines 3, 4).

Auxiliary task. When annotations are available, our learning tasks can be used as an auxiliary to the primary loss. Using the 3DNotes-SN’s annotated training set to train the primary supervised task, we set the *comb* task as its auxiliary counterpart by summing the losses. We evaluate zero-shot performance on COMETA-ZS. We observe an 8% increase in R@25, illustrated in Fig. 2. Since most annotations cover no more than a couple thousands concepts, which is a tiny portion of a typical medical KG’s size, this demonstrates the generalizing capacity of our approach on the vast majority of unseen concepts.

Private KG. In practice, if the target medical ontology is a private KG (Wise et al., 2020; Bhatia et al., 2020), one can also consider customizing the learning tasks that follow the synonym and graph-based frameworks outlined in this work to bring greater gains.

7 Conclusion

We present a framework for allowing entity retrieval models to mine rich semantics from a medical KG. We show its effectiveness in zero-shot set-

tings through extensive experiments. In addition, we demonstrate the ease with which the framework can be adapted to serve as an auxiliary task when annotations are available. Future research should explore more fine-grained approaches to combine tasks.

References

- Khushbu Agarwal, Tome Eftimov, Raghavendra Ad-danki, Sutana Choudhury, Suzanne Tamang, and Robert Rallo. 2019. Snomed2vec: Poincare and random walk embeddings of a clinical knowledge base for healthcare analytics. In *KDD Workshop on Applied Data Science for Healthcare: Bridging the Gap between Data and Knowledge*.
- Emily Alsentzer, John R. Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew B. A. McDermott. 2019. [Publicly available clinical bert embeddings](#).
- Marco Basaldella, Fangyu Liu, Ehsan Shareghi, and Nigel Collier. 2020. [COMETA: A corpus for medical entity linking in the social media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3122–3137, Melbourne, Australia. Association for Computational Linguistics.
- Parminder Bhatia, Lan Liu, Kristjan Arumae, Nima Pourdamghani, Suyog Deshpande, Ben Snively, Mona Mona, Colby Wise, George Price, Shyam Ramaswamy, Xiaofei Ma, Ramesh Nallapati, Zhiheng Huang, Bing Xiang, and Taha Kass-Hout. 2020. [Aws cord-19 search: A neural search engine for covid-19 literature](#).
- Wei-Cheng Chang, Felix X. Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. [Pre-training tasks for embedding-based large-scale retrieval](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Daniel Hanisch, Katrin Fundel, Heinz-Theodor Mevis-sen, Ralf Zimmer, and Juliane Fluck. 2005. [Prominer: rule-based protein and gene entity recognition](#). *BMC Bioinformatics*, 6:1:S14.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*.
- Sunil Mohan and Donghui Li. 2019. [Medmentions: A large biomedical corpus annotated with UMLS concepts](#). *CoRR*, abs/1902.09476.
- Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2011. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Qiong Wang, Zongcheng Ji, Jingqi Wang, Stephen Wu, Weiyan Lin, Wenzhen Li, Li Ke, Guohong Xiao, Qing Jiang, Hua Xu, and Yi Zhou. 2020. [A study of entity-linking methods for normalizing chinese diagnosis and procedure terms to icd codes](#). *Journal of Biomedical Informatics*, 105:103418.
- Colby Wise, Vassilis N. Ioannidis, Miguel Romero Calvo, Xiang Song, George Price, Ninad Kulkarni, Ryan Brand, Parminder Bhatia, and George Karypis. 2020. [Covid-19 knowledge graph: Accelerating information retrieval and discovery for scientific literature](#).
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. [Zero-shot entity linking with dense entity retrieval](#). In *EMNLP*.
- Ming Zhu, Busra Celikkaya, Parminder Bhatia, and Chandan K. Reddy. 2020. [Latte: Latent type modeling for biomedical entity linking](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9757–9764.