



# Reveal the Unknown: Out-of-Knowledge-Base Mention Discovery with Entity Linking

Hang Dong  
University of Oxford  
hang.dong@cs.ox.ac.uk

Jiaoyan Chen  
The University of Manchester &  
University of Oxford  
jiaoyan.chen@manchester.ac.uk

Yuan He  
University of Oxford  
yuan.he@cs.ox.ac.uk

Yinan Liu  
Nankai University  
liuyn@mail.nankai.edu.cn

Ian Horrocks  
University of Oxford  
ian.horrocks@cs.ox.ac.uk

## ABSTRACT

Discovering entity mentions that are out of a Knowledge Base (KB) from texts plays a critical role in KB maintenance, but has not yet been fully explored. The current methods are mostly limited to the simple threshold-based approach and feature-based classification, and the datasets for evaluation are relatively rare. We propose BLINKout, a new BERT-based Entity Linking (EL) method which can identify mentions that do not have corresponding KB entities by matching them to a special NIL entity. To better utilize BERT, we propose new techniques including NIL entity representation and classification, with synonym enhancement. We also apply KB Pruning and Versioning strategies to automatically construct out-of-KB datasets from common in-KB EL datasets. Results on five datasets of clinical notes, biomedical publications, and Wikipedia articles in various domains show the advantages of BLINKout over existing methods to identify out-of-KB mentions for the medical ontologies, UMLS, SNOMED CT, and the general KB, WikiData.<sup>1</sup>

## CCS CONCEPTS

• **Computing methodologies** → **Information extraction**; *Ontology engineering*.

## KEYWORDS

Entity Linking, Knowledge Base Enrichment, Language Models, Biomedical Ontologies, WikiData

### ACM Reference Format:

Hang Dong, Jiaoyan Chen, Yuan He, Yinan Liu, and Ian Horrocks. 2023. Reveal the Unknown: Out-of-Knowledge-Base Mention Discovery with Entity Linking. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3615036>

<sup>1</sup>The implementation, dataset construction scripts, and datasets are publicly available at <https://github.com/KRR-Oxford/BLINKout> and <https://zenodo.org/record/8228371>.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0124-5/23/10.

<https://doi.org/10.1145/3583780.3615036>

## 1 INTRODUCTION

Knowledge Bases (KBs) are widely used for representing entities and facts about the world with reasoning supported. KBs are inherently incomplete. New entities are constantly emerging, for example, from late 2020 to early 2022, a new variant of SARS-CoV-2 emerged every few months [44]. Existing KBs may thus also inevitably miss entities, for example, “Curry-Jones syndrome” [42] was not added to SNOMED CT [9] until 2017. Delays in incorporating these entities into the KB may result in failure to cataloguing, searching, and reasoning with them. Automated discovery of mentions of new entities from common resources such as texts can support the maintenance of KBs.

Updating KBs with entities from texts is highly relevant to *Entity Linking* (EL), which is to match mentions in texts to entities in a KB [34, 40]. Current works on EL, however, often assume all the target mentions have corresponding entities in the KB and ignore mentions that have no corresponding entities [3, 45]. The latter are sometimes called *out-of-KB mentions* and are matched with a *NIL entity*. To maintain a KB, it is required to discover out-of-KB mentions which can be further processed as new KB entities.<sup>2</sup>

As NIL is not described in the KB, it is hard to obtain their lexical or embedding representations, and consequently, it is hard to predict them directly. One idea is setting a threshold towards the mention-to-entity matching score: a mention is regarded as NIL if its scores to all the KB entities are below the threshold [6]. Another idea is to *classify* a mention into NIL or in-KB based on a set of features regarding the mention and its top-*k* entity candidates [46]. There are also some other studies attempting to represent the NIL entity with key phrases or features based on external corpora [21] or manual effort [22].

Recently, neural EL has shown good performance by applying pre-trained language models (LMs), e.g., BERT [8], to represent texts and entities [45]. In contrast, such LM-based methods have been rarely applied for out-of-KB mentions, as reviewed in [38, 39]. Traditional, threshold-based approaches in combination with deep learning are still the state-of-the-art [7, 25]. There is a lack of studies to seamlessly integrate out-of-KB mention discovery with pre-trained LM-based approaches like BLINK [45]. Also, entities are likely to have various surface forms or synonyms [7, 11]. This entity variety suggests enhancing synonyms for EL and may help differentiate between in-KB and out-of-KB entities.

<sup>2</sup>In this work, we focus on identifying out-of-KB *mentions*, and leave the canonicalisation and placement of the new entities in the KB for future work.

Besides the shortage of methodology research, benchmarking datasets considering out-of-KB mention discovery are also relatively rare compared to in-KB EL. Most EL datasets assume that the KB is complete and do not include out-of-KB entities, e.g., the MedMentions dataset [32] that links mentions to concepts in UMLS [5]. The most recent large-scale dataset is NILK [23], which synthesizes NIL entities from the entity gap between the newer version of Wikidata (2021) and the older version (2017) to enrich the older KB. Also, in the biomedical domain, the main out-of-KB EL dataset is Share/CLEF 2013 [41], with NIL mentions manually identified for training and evaluation. The only two other datasets from the other domains also rely on manual annotations, e.g., historical newspapers [12] and news in microposts [37], which require domain experts' substantial effort. Also, all the previous studies only focus on a single strategy (either manual effort or KB versioning) to construct EL datasets with NIL labels.

In this work, (i) we define the out-of-KB mention discovery problem and propose a method named BLINKout, based on the BERT-based EL method, BLINK [45], where new techniques including NIL entity representation & classification and synonym enhancement are developed and applied; (ii) we summarize and apply strategies to automatically construct out-of-KB mention discovery benchmarks from an in-KB EL dataset by pruning or using older versions of the linked ontology, besides manual labelling. Five out-of-KB mention discovery datasets are selected and constructed using three data strategies (i.e., manual labelling, KB pruning, KB versioning), with texts of clinical notes, biomedical publications and Wikipedia articles, two medical ontologies (UMLS and SNOMED CT), and one general purpose KB (WikiData) covering various domains.

Experimental results on the datasets show the advantage of the BLINKout approach for out-of-KB mention discovery, with comparison to rule-based, threshold-based, and feature-based baselines and ablation studies, up to nearly 40% improvement of out-of-KB  $F_1$  score compared to the second best baseline.

## 2 PRELIMINARIES

### 2.1 Problem Definition

Given a KB containing a set of entities  $E$  and a text corpus in which a set of mentions  $M$  are identified in advance, *Entity Linking* (EL, a.k.a. Entity Disambiguation or Entity Normalisation) is to map each mention  $m$  in  $M$  to its corresponding entity  $e$  in  $E$  [40]. Due to the limited knowledge coverage of the KB, it is possible that there is no entity in  $E$  that can be matched with a given mention. Thus we have the following more general problem of **out-of-KB mention discovery**, formulated as an extended EL task that focuses on identifying out-of-KB mentions along with linking in-KB mentions.

Given a text corpus with a set of identified mentions (each in a context window in a document)  $D_M$  and a KB with entities  $E$ , *out-of-KB mention discovery* is to develop a function  $f$  such that each mention  $d_m$  in  $D_M$  is mapped to an item  $f(d_m) \in E \cup \{\text{NIL}\}$ , where NIL is a special entity denoting that there are no entities in  $E$  that can be matched with  $d_m$ . In this study, we focus on both ontologies (e.g., SNOMED-CT) and general KBs (e.g., WikiData). Ontologies are a common type of KB that is often defined as a shared, explicit specification of a conceptualisation of a domain [15], and we consider an ontology's classes (or concepts) as its

entities for the linking. Note that each entity may have definitions and synonyms that can be utilized for out-of-KB mention discovery.

### 2.2 BERT-based Entity Linking

We summarize the BERT-based Entity Linking methods (e.g., BLINK [45]) below. They usually have two stages: *candidate creation* and *candidate ranking* [25, 34, 45]. In candidate creation, the approaches aim to narrow down the vast number of entities into a manageable subset (e.g., tens or hundreds), and in candidate ranking, the approaches aim to rank the candidate entities of each mention according to the probability that they match the given mention.

**Candidate Creation with Bi-encoder.** The bi-encoder fine-tunes two BERT models  $BERT_m$  and  $BERT_e$  to embed mentions and entities, resp., into a dense embedding space, so that given a new mention, the nearest candidates can be easily retrieved. For a mention  $d_m$  and an entity  $d_e$ , their embeddings can be accessed as

$$v_m = \text{red}(BERT_m(d_m)); \quad v_e = \text{red}(BERT_e(d_e)) \quad (1)$$

where  $\text{red}(\cdot)$  denotes a function for extracting the vector representation of a textual sequence from a BERT model using its last layer representation of the [CLS] token.

Mentions and entities can be fed into the BERT models in different ways. In the classic work [45], a mention is fed into  $BERT_m$  as [CLS]  $\text{ctx}_l$  [M<sub>s</sub>] mention [M<sub>e</sub>]  $\text{ctx}_r$  [SEP], where  $\text{ctx}_l$  and  $\text{ctx}_r$  are the left and right contexts of the mention in the document, resp.; and an entity (with name and definition) is fed into  $BERT_e$  as [CLS] name [ENT] definition [SEP]. Note that [M<sub>s</sub>], [M<sub>e</sub>], and [ENT] are special tokens for separation.

With these embeddings, the **score** which measures the similarity between a mention and an entity can be calculated, e.g., by their dot-product defined as  $s(m, e) = v_m \cdot v_e$ . The scores are further used to generate top- $k$  candidates for each mention.

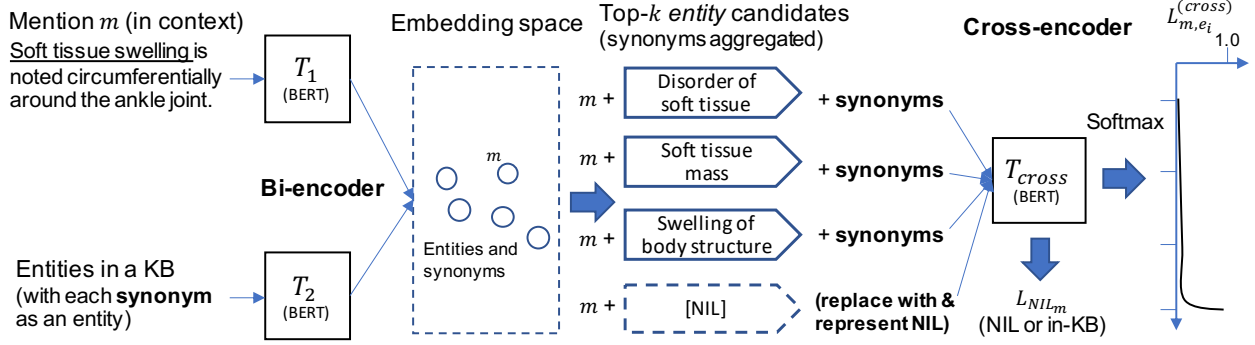
Bi-encoder can be trained with a loss function to make each mention close to its matched entity in the embedding space, but far away from the other entities within the same batch. This can be realised with the max-margin triplet loss [3, 35] described below<sup>3</sup>, where  $\alpha$  is a margin of small value (e.g., 0.2) and  $[x]_+$  denotes  $\max(x, 0)$ , for each mention to its gold entity (the  $i$ -th) in a batch.

$$L_{m_i, e_i} = \sum_{j \neq i} [\alpha - s(m_i, e_i) + s(m_i, e_j)]_+ \quad (2)$$

**Candidate Ranking with Cross-encoder.** Given that the bi-encoder can only coarsely identify top- $k$  candidates but may not correctly rank them among a large number of KB entities [45], we use the cross-encoder to make the candidate ranking. The cross-encoder is fed with information of both a mention and its top- $k$  entities from the bi-encoder, and performs a multi-class classification, i.e., the entity that is more likely to be matched when the mention is predicted with a higher score.

Each mention  $m$  is concatenated with each of its top- $k$  entities following the same input format as in the bi-encoder. The concatenation of the input for the mention and the entity  $e$  (without the [SEP] and [CLS] tokens in between) is denoted as  $d_{m,e}$ , i.e., [CLS]

<sup>3</sup>The bi-encoder is originally trained using a cross-entropy loss [45] (or a non-temperated contrastive loss cf. Gao et al. [14]) for a similar intuition. We found that the triplet loss led to a more stable convergence than the cross-entropy loss in [45], with a similar performance to the latter, in the experiments.



**Figure 1: BLINKout architecture for out-of-KB mention discovery, adapting BERT-based Entity Linking [45]: bi-encoder encodes separately the mention  $m$  in a context and the entities (each synonym as an entity) into a dense embedding space; cross-encoder classifies the most relevant entity candidate (with synonym concatenated), with NIL Entity Representation & Classification that appends a [NIL] special token to replace the last candidate (if NIL was not predicted by the bi-encoder), jointly learned with  $L_{NIL}$  to classify whether the mention is out-of-KB.**

$ctx_t_l [M_s]$  mention  $[M_e]$   $ctx_t_r [SEP]$  name [ENT] definition [SEP]. The input  $d_{m,e}$  is fed into the cross-encoder, which is composed of a BERT model and a linear layer, for a score output  $s_{m,e}^{(cross)}$ . The vector  $w$  constitutes the parameters to be learned in the model.

$$s_{m,e}^{(cross)} = v_{m,e}^{(cross)} w; \quad v_{m,e}^{(cross)} = \text{red}(BERT_{cross}(d_{m,e})) \quad (3)$$

The cross-encoder is learned with the following cross-entropy loss with softmax activation, where  $E_{cand}$  denotes the indices of the top- $k$  candidate entities of  $m$ , and  $i$  is the index of the gold entity.

$$L_{m,ei}^{(cross)} = -\log\left(\frac{\exp(s_{m,ei}^{(cross)})}{\sum_{j \in E_{cand}} \exp(s_{m,e_j}^{(cross)})}\right) \quad (4)$$

### 3 BLINKOUT

The BERT-based EL [45] described previously in Section 2.2 does not consider the NIL entity. We present NIL entity representation & classification, an approach for out-of-KB mention discovery that adapts the cross-encoder, applicable either with or without training, and discuss the techniques to represent the NIL entity with LMs. Then, we propose methods to enhance both bi-encoder and cross-encoder with synonyms or variants which are prevalent for entities.

#### 3.1 NIL Entity Classification

For the classification approach, we ensure that the NIL entity is within the top- $k$  candidates by replacing the last candidate with NIL. This gives a chance for the cross-encoder with  $L^{(cross)}$  to potentially classify NIL as the top entity for the mention, as inspired by [34]. The textual representation of NIL, encoded with BERT-like LMs, will affect the performance of classification; we discuss different NIL entity representations in 3.2.

**Joint Training for NIL Classification.** We add a joint loss to learn to verify whether a mention is NIL or in-KB, conditioned on the mention representation, similar to [16]. The sigmoid function  $\sigma$  with a linear layer is used to form a probability score,  $s_{isNIL_m}$ , given the mention representation. The vector  $w_{NIL}$  constitutes the model parameters, learned with the binary cross-entropy loss below.

$$s_{isNIL_m} = \sigma(v_m^{(cross)} w_{NIL}); \quad v_m^{(cross)} = \text{red}(BERT_{cross}(d_m)) \quad (5)$$

$$L_{NIL_m} = -y_i \log(s_{isNIL_m}) + (1 - y_i) \log(1 - s_{isNIL_m}) \quad (6)$$

The overall loss with joint training for the cross-encoder is described below. The best value of  $\lambda_{NIL}$  varied from 0.01 to 0.25 across the datasets based on parameter tuning with the validation set.

$$L_m = L_{m,ei}^{(cross)} + \lambda_{NIL} * L_{NIL_m} \quad (7)$$

#### 3.2 NIL Entity Representation

The NIL entity does not have a native textual representation in the KB. A better representation of NIL will help the LM to discriminate NIL from in-KB entities. This representation can be either static, fixed (or unsupervised) or dynamic, fine-tuned (or supervised with NIL mentions) in the LM.

We represent NIL as a special token [NIL] by taking advantage of the tokenizer in a BERT-like LM. This assigns special semantics to the NIL entity so it is not confused with the names of other entities in the KB. Also, the continuous representation of [NIL] can be further fine-tuned with the LM. A more naive representation is "NIL", with the definition of "It is a NIL option.", used in a previous study [16]; we also replace NIL with [NIL] in the definition. Similar to in-KB entities, we add the [ENT] special token between the name and the definition. A list of NIL entity representations and their results are presented in Table 3. Our final approach uses the dynamic, fine-tuned [NIL] representation that leverages the labelled out-of-KB mentions in the training data.

#### 3.3 Synonym Enhancement

The original BLINK [45] focuses on Wikipedia texts and does not consider synonyms. Synonyms are prevalent for real-world entities (see data statistics in Table 1), e.g., entity C0428977 in the UMLS has the name of "Bradycardia" and some synonyms such as "Slow heart beat" and "Heart rare slow". When in-KB entities are better represented with synonyms, the out-of-KB mentions can be more precisely identified by the LM. We thus enhance the bi-encoder and the cross-encoder with synonyms, with two different approaches,

resp.: (i) augmentation of each synonym as an entity in the bi-encoder, for its thorough training, and (ii) concatenation with [SYN] special token in the cross-encoder, for its efficient training.

**Synonym Augmentation in Bi-encoder.** We use synonyms for data augmentation to enhance the training. Each synonym is treated as a separate entity to be matched to the mention. This can significantly augment the size of the training data. After the scoring, we aggregate the entities and synonyms into top- $k$  unique entity candidates by setting the score of an entity as its highest score among all its variations.

**Synonym Concatenation in Cross-encoder.** It is inefficient and infeasible to use synonym augmentation in the cross-encoder, as the number of classes can be significantly increased (e.g., by around 3-4 times for the UMLS and SNOMED CT subset) and unstable or non-fixed. Thus, we model each entity candidate  $e$  with the concatenation of its synonyms, separated by the [SYN] special tokens, i.e., [CLS] name [ENT] synonym\_1 [SYN] synonym\_2 . . . synonym\_n [SYN] definition.<sup>4</sup> This keeps the number of entities to classify to  $k$  instead of treating each synonym as an entity, and thus significantly reduces the computation while still making full use of the synonyms.

Finally, we use **BLINKout** to refer to the approach that integrates BERT-based EL with synonym enhancement and fine-tuned NIL entity representation for classification.

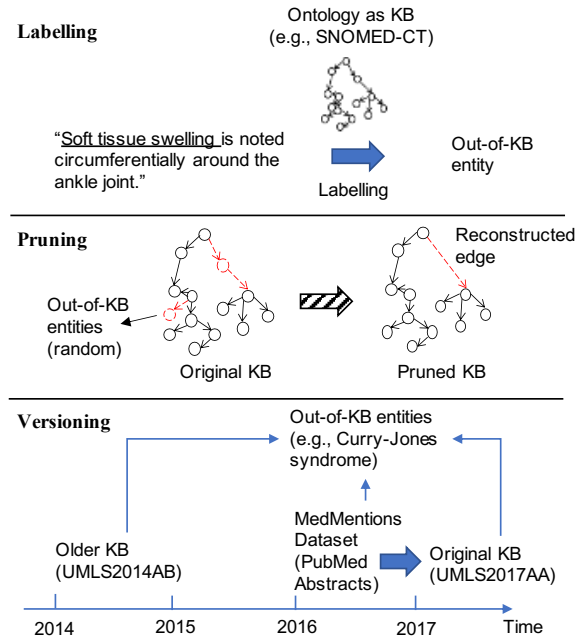
## 4 DATASET CONSTRUCTION

There are different strategies to construct NIL-enhanced (or NIL-labelled) EL datasets, which contain out-of-KB mentions labelled with NIL (i.e., each mention is linked to either an entity in  $E$  or NIL). One straightforward way is **Manual Labelling**. We adopt one main manually NIL-labelled dataset in the medical domain, ShARe/CLEF 2013, which consists of clinical notes (discharge summaries and electrocardiogram, echocardiogram, and radiology reports) in the version 2.5 of the MIMIC-II dataset [41]. The target ontology is SNOMED CT (represented with their mapped UMLS CUI) refined with the Disorder semantic group of 10 semantic types in the UMLS, defined in the annotation guideline [13]. We use UMLS 2012AB following Ji et al. [25]. Around 30% of the mentions are out-of-KB.<sup>5</sup>

Manual labelling costs much labour. We thus apply two automatic strategies, KB Pruning and KB Versioning, to synthesize out-of-KB mentions within normal EL datasets.

**KB Pruning.** We randomly sample a portion (e.g., 10% and 20%) of the entities in the target KB and remove them from the KB. To preserve the hierarchies (formed by subsumption relations, a.k.a. “subclass of” or “isA” relations) in an ontology, we link the parent and child of each removed entity as in He et al. [18, 19]<sup>6</sup> and this forms a new ontology. Mentions that are originally linked to the removed entities are labelled by NIL.

**KB Versioning.** We also consider replacing the current version of the KB with an older version, so that mentions linked to entities



**Figure 2: Strategies for NIL-enhanced EL dataset construction: Manual Labelling, KB Pruning, and KB Versioning. The specified UMLS ontologies are the subset linked to SNOMED CT under the semantic type of T047, “Disease or Syndrome”.**

which are in the newer version but not in the older version become out-of-KB. Using the MedMentions dataset [32] as an example (see Figure 2), UMLS2017AA was used to exhaustively annotate entities in publication titles and abstracts in 2016; we chose its older version of UMLS2014AB (released in Nov 2014). We narrowed the UMLS semantic type to T047, *Disease or Syndrome*, and selected the source as “SNOMEDCT\_US” (similar to ShARe/CLEF 2013). We gathered the entities that are in UMLS2017AA but not in the earlier version<sup>7</sup>. Mentions labelled with these new entities become out-of-KB<sup>8</sup>.

We also adapted a recent general domain dataset constructed using the KB versioning approach. The recent NILK dataset [23] is created on the general, multi-domain KB, WikiData [43]. The WikiData\_2021 version is used to supplement entities in WikiData\_2017. In-text hyperlinks (each link to another Wikipedia article) in the Wikipedia 2017 dump<sup>9</sup> are used as the source of mentions; the hyperlinked Wikipedia articles are mapped to WikiData ID to generate in-KB (in WikiData\_2017) and out-of-KB entities (not in WikiData\_2017 but in WikiData\_2021) for the mentions [23]. We benchmark with a proportional random sample of 0.1% of the mentions for each of the original data splits, given that the whole NILK dataset is huge (over 107M mentions each to be linked to one of

<sup>4</sup>For bi-encoder, it is also possible to use synonym concatenation, while the proposed synonym augmentation performed better on Share/CLEF 2013 in the experiments.

<sup>5</sup>The out-of-KB mentions can be out of the Disorder semantic group but belong to another semantic type (e.g., “[Thick] bronchial secretions” only as body substance; “soft tissue swelling” and “vomiting coffee grounds” only as findings) [13].

<sup>6</sup>We used the pruning function implemented in the Python library, DeepOnto [18], see <https://krr-oxford.github.io/DeepOnto/deeponto/onto/pruning/>.

<sup>7</sup>We further filtered out those that were merged to a concept in the previous version of the ontology (as defined in Retired CUI Mapping, MRCULRRF in the UMLS), representing around 6% of all newly added concepts.

<sup>8</sup>Examples include “CJS” or “Curry-Jones syndrome” (C0795915), “CKD” or “Chronic Kidney Diseases” (C1561643), and “Pandemic influenza” (C4304383), which were added to the UMLS between Nov 2014 and May 2017 under the semantic type T047, Disease or Syndrome.

<sup>9</sup><https://archive.org/download/enwiki-20170220/enwiki-20170220-pages-articles.xml.bz2>.

	ShARe/CLEF 2013 train / test	MM-pruned-0.1 train / valid / test	MM-pruned-0.2 train / valid / test	MM-2014AB train / valid / test	NILK-sample train / valid / test
Out-of-KB mention creation	Manual Labelling	KB Pruning	KB Pruning	KB Versioning	KB Versioning
Domain or document type	Clinical notes	PubMed abstracts	PubMed abstracts	PubMed abstracts	Wikipedia articles
# of docs	199 / 99	2,635 / 878 / 879	2,635 / 878 / 879	2,635 / 878 / 879	81,334 / 10,587 / 10,528
# of mentions	5,816 / 5,351	6,201 / 2,121 / 2,000	6,201 / 2,121 / 2,000	6,181 / 2,112 / 1,988	86,379 / 10,688 / 10,613
# of out-of-KB mentions	1,639 / 1,723	456 / 201 / 155	1,031 / 405 / 415	307 / 82 / 103	1,312 / 154 / 158
% of out-of-KB mentions	28.2% / 32.3%	7.3% / 9.5% / 7.8%	16.6% / 19.1% / 20.8%	5.0% / 3.9% / 5.2%	1.5% / 1.4% / 1.5%
# of in-KB mentions	4,109 / 3,604	5,745 / 1,920 / 1,845	5,170 / 1,716 / 1,585	5,874 / 2,030 / 1,885	85,067 / 10,534 / 10,455
% mentions w/ zero-shot entities	- / 11.6%	- / 24.8% / 25.3%	- / 23.2% / 22.0%	- / 27.8% / 28.0%	- / 98.6% / 98.5%
# of entities in KB	88,150	35,392	31,460	35,398	79,412
# of entities & synonyms in KB	288,490	126,188	112,097	124,132	121,191
% of entities having synonyms	99.0%	99.4%	99.4%	99.3%	25.5%
Ave # syns per entity (having syns)	2.3 (2.3)	2.6 (2.6)	2.6 (2.6)	2.5 (2.5)	0.5 (2.1)

**Table 1: Statistics of out-of-KB linking datasets. Slashes separate the statistics of training, validation, and testing sets. “MM” denotes MedMentions; pruned-0.1 and pruned-0.2 denote the percentage of pruned concepts in the ontology; “2014AB” refers to the older version of the UMLS applied; “NILK-sample” denotes the random sample of the NILK dataset that enriches WikiData.**

14.6M English-described entities in the full WikiData dump). We also select entities in the WikiData KB which have at least one mention linked after the random sampling of the mentions. This results in data sample and entity sizes comparable to the other four datasets and keeps the same proportion of NIL mentions as the original data: around 1.5%, much lower than the other datasets.

The advantage of KB Pruning is that it allows us to simulate the creation of out-of-KB entities and control their percentage, while KB Versioning creates datasets that are closer to real-life situations where new concepts emerge over time.

Statistics of the datasets are shown in Table 1, including the counts of mentions and documents, and the percentages of out-of-KB mentions, zero-shot mentions, and entities with synonyms. The KB Versioning approach resulted in a much lower percentage of out-of-KB mentions. Also, NILK-sample is a zero-shot EL dataset (from its original construction in [23]), i.e., no overlapping entities (except NIL) among the training, validation, and testing sets, while in the other datasets most of the entities are seen in training. We also note that only about a quarter (25.5%) of the sampled WikiData entities are associated with at least one synonym, compared to nearly all (>99.0%) entities in the subset of UMLS and SNOMED-CT having synonym(s). The average number of synonyms per entity (having at least one synonym) in the NILK-sample dataset is also lower, compared to the other datasets.

## 5 EVALUATION

### 5.1 Comparison Methods

The main baseline methods for out-of-KB mention discovery are:

- (i) **Rule-based approach** (or Sieve-based): using carefully crafted rules (each as a “sieve”) designed for biomedical texts, where NIL is detected when no in-KB entity can be linked to [11].
- (ii) **Threshold-based approach**: setting a threshold on the EL system’s prediction score for each candidate, if the score for the closest in-KB candidate is still below the threshold, the mention is out-of-KB [7, 24]; this is applied with BM25 for candidate generation and BERT-based cross-encoder for candidate ranking and synonym enhancement (“Th-based BM25+BERT+syn”) [25], where we tuned the threshold (as 0.85) and the domain-specific setting, as SapBERT [27] (“Th-based BM25+SapBERT+syn”), based on the validation set.
- (iii) **Feature-based NIL entity classification** (“Ft-based Classifier”): including

string-matching features [10, 31, 34], entity contextual feature space, and embedding-based feature space [46]; we also provide a feature-based baseline that uses the entity candidates and dynamic features (i.e.,  $v_{m,e}^{(cross)}$ ) in BLINKout (“Ft-based BLINKout  $L_{NIL}^{(Ft)}$ ”).

For model comparison, the default number of top- $k$  was set as 10 for the baselines using BM25, BLINK, and BLINKout, following Ji et al. [25]. In the “ $k$  and NIL rep tuned” setting, both  $k$  and NIL entity representation were tuned based on the validation set to optimise  $F_{1o}$  (see Section 5.4 and Appendix B for details in parameter tuning, with the optimal settings in Appendix A).

**Adapting BLINK without Training.** We also implemented both threshold-based and the NIL entity representation approach to a BLINK model trained from all in-KB entities (“Th-based BLINK”<sup>10</sup> and “NIL-rep-based BLINK”), either or not using synonyms (“+syn”).

### 5.2 Evaluation Metrics

In the conventional EL setting, which does not discriminate between in-KB and out-of-KB mentions, the *overall* (or in-KB+out-of-KB) accuracy ( $A$ ), precision ( $P$ ), recall ( $R$ ), and  $F_1$  scores are the same, i.e.,  $A = P = R = F_1 = \frac{TP}{|M|}$ , where  $TP$  and  $|M|$  denote the number of correctly linked mentions and all the target mentions, resp. [38, 40].

In this work, we propose to use *out-of-KB* precision ( $P_o$ ), recall ( $R_o$ ), and  $F_1$  ( $F_{1o}$ ) scores to measure how well out-of-KB mentions are detected. Analogously, we can calculate the *in-KB* precision ( $P_{in}$ ), recall ( $R_{in}$ ), and  $F_{1in}$  scores.

$$P_o = \frac{TP_o}{TP_o + FP_o}; R_o = \frac{TP_o}{TP_o + FN_o} \quad (8)$$

$$P_{in} = \frac{TP_{in}}{TP_{in} + FP_{in}}; R_{in} = \frac{TP_{in}}{TP_{in} + FN_{in}} \quad (9)$$

where  $TP_o$ ,  $FP_o$ , and  $FN_o$  (resp.,  $TP_{in}$ ,  $FP_{in}$ , and  $FN_{in}$ ) are the numbers of true positive, false positive, and false negative out-of-KB (resp., in-KB) mentions.  $TP_{in}$  refers to the number of mentions that are predicted with the correct entities in the KB instead of mentions that are simply predicted as in-KB.  $F_1$  is the harmonic mean of  $P$  and  $R$ .

<sup>10</sup>The out-of-KB threshold is placed at the cross-encoder. We also ensure that the NIL entity is in the bi-encoder’s top- $k$  candidates by replacing the  $k$ -th candidate to NIL, thus adding an extra threshold for the bi-encoder does not affect the results.



### 5.3 Results

**Out-of-KB Metrics.** Table 2 shows the main results for out-of-KB mention discovery. BLINKout, with SapBERT and for NILK-sample with BERT, achieved the top out-of-KB  $F_1$  ( $F_{1_o}$ ) on all datasets.

The rule-based method (Sieve-based) [11] was effective on biomedical datasets, given the regular syntactic structure of the biomedical texts and concepts, but less effective on the general domain (NILK-sample). The rule-based approach resulted in high out-of-KB recall scores, but still lacked in out-of-KB precision and  $F_{1_o}$  compared to BLINKout; this is because the former tends to result in a large number of out-of-KB mentions using the “NIL as no candidate” rule.

Feature-based methods, facing the challenge of imbalanced binary classification, also tend to predict a larger number of NILs and result in a high out-of-KB recall, but low precision and  $F_1$  scores.

Threshold-based approaches provide higher out-of-KB precision than feature-based, but lower metrics than the rule-based approach. The setting using synonyms with bi-encoder (“BLINK+syn”) instead of BM25 (“BM25+BERT+syn”) greatly improved the in-KB results and thus the out-of-KB results based on a threshold.

NIL-rep-based approaches with fixed NIL entity representations performed better than or on par with threshold-based approaches. The special token [NIL], even fixed (unsupervised), can discriminate NIL from in-KB entities using the weights in pre-trained LM.

BLINKout outperformed all the baselines on all the datasets. A large margin was achieved compared to the rule-based (second best for biomedical datasets, i.e., except for NILK-sample): for MM-2014AB and NILK-sample, by about 30-40% improvement of  $F_{1_o}$ .

For NILK-sample, the improvement with synonyms was less obvious. This is due to the zero-shot nature of the dataset and inadequacy of synonyms: first, NILK is a zero-shot entity linking dataset, while synonyms most benefit seen entities rather than unseen entities<sup>11</sup>; second, the percentage of entities having a synonym is much lower in NILK-sample (25%) than the other datasets (99%) and the average number of synonyms per entity is lower.

**Ablation Study.** The contributing components in BLINKout are: (i) **fine-tuned NIL entity representation** with NIL-labelled training data, e.g., with 3-18% improvement of  $F_{1_o}$  between “BLINKout (BERT)” and “NIL-rep-based BLINK+syn”; (ii) **synonym enhancement**, e.g., with 2-14% improvement of  $F_{1_o}$  between “BLINKout (BERT)” and “BLINKout (-syn, BERT)”, except for a marginal improvement (of 0.4%) on the NILK-sample dataset; (iii) **domain-specific LMs**, e.g., with 3-19% improvement of  $F_{1_o}$  between “BLINKout (SapBERT)” and “BLINKout (BERT)”, except for a marginal improvement (of 0.7%) on the NILK-sample dataset.

**Overall Accuracy.** For the overall accuracy on both in-KB and out-of-KB mentions ( $A$ ), the proposed models perform the best or competitively in all datasets. This shows that the approaches to discover out-of-KB mentions do not compromise the performance of in-KB entities. In-KB EL results are in Table 5 in Appendix C.

### 5.4 Sensitive Settings

The results of BLINKout are sensitive to the settings below.

<sup>11</sup>We also tested the zero-shot (ZS) version of the other four datasets (by removing testing mentions linked to overlapped entities to the training set), and found that using synonyms may not improve much (compared to the original data setting) or may even impede the performance of linking, see results in Appendix D.

**The Number of Top- $k$  Candidates.** Out-of-KB metric scores generally have a positive correlation with  $k$  when it is below a limit. A higher top- $k$  may make it harder to train a model to discriminate out-of-KB from the remaining classes using the cross-encoder; this especially decreases the performance when there are limited out-of-KB samples for training (e.g., in MM-2014AB). A lower top- $k$  is likely to decrease in-KB EL results. Details are in Appendix B.

**NIL Entity Representation.** NIL representation sets a prior anchor point in the embedding space for the bi-encoder and is used in the cross-encoder for candidate ranking. Results on ShARe/CLEF 2013 and NILK-sample are displayed in Table 3: [NIL] representation performed the best in most settings, either fixed or fine-tuned.

**Domain-Specific LMs.** In-domain and knowledge-enhanced LMs, SapBERT (and also PubMedBERT [17]), obtained better in- and out-of- KB results for datasets in the biomedical domain.

## 6 QUALITATIVE ANALYSIS

We selected samples from the test set in the datasets regarding erroneous predictions for out-of-KB mentions, displayed in Table 4. NILK-sample dataset has many text sequences rendered from Tables (see rows 5-6), which provides a different context compared to other datasets, and fine-grained entities such as names of cities in other languages, asteroids, and animated series (see rows 5-7).

BLINKout models can correctly identify NIL mentions, while false but similar in-KB entities were predicted in *Th*- and/or NIL-rep- based methods (see rows 3, 5-6). BLINKout with joint learning affects the prediction score for NIL, which can result in both true or false positive NIL predictions (*cf.* row 1 and 2). The NIL-rep-based approach can sensitively identify NIL mentions with a high rank (if not predicted as the top, in rows 1, 3-7), without training with NIL-labelled mentions. The threshold-based approach may likely result in false negative predictions for out-of-KB mentions when the incorrect but similar entity was predicted with a very high score near 100% (as in rows 1, 4, 6-7). Finally, there are challenging out-of-KB entities which are very similar to in-KB entities in the KB pruning based dataset, MM-pruned-0.1, and fine-grained entities in NILK-sample, so that all models yielded a wrong prediction, but with a high rank of NIL as the 2nd or 3rd (in rows 4, 7).

## 7 DISCUSSION

**Results w.r.t. Dataset Construction Strategies.** The out-of-KB performance of BLINKout was the highest on the Manual Labelling data, ShARe/CLEF 2013, and lowest on the KB Versioning data, NILK-sample. This is partially related to the percentage of out-of-KB mentions: the lower the percentage in the training data, the more challenging their detection. On the more realistic KB Versioning dataset, BLINKout outperforms the rule-based approach with a larger gap compared to datasets created with the other strategies.

**NIL Entity Representation.** The results show that NIL entity representation based methods perform better or on par with the threshold-based method, and also have a high rank of NIL among in-KB entities for out-of-KB mentions. BLINKout with the fine-tuned [NIL] representation yielded a large margin of improvement. It is safer to use non-natural language tokens, i.e., [NIL], rather than “NIL” to avoid confusion with other vocabularies.

	ShARe/CLEF 2013				MM-pruned-0.1				MM-pruned-0.2				MM-2014AB				NILK-sample			
	A	P <sub>o</sub>	R <sub>o</sub>	F <sub>1<sub>o</sub></sub>	A	P <sub>o</sub>	R <sub>o</sub>	F <sub>1<sub>o</sub></sub>	A	P <sub>o</sub>	R <sub>o</sub>	F <sub>1<sub>o</sub></sub>	A	P <sub>o</sub>	R <sub>o</sub>	F <sub>1<sub>o</sub></sub>	A	P <sub>o</sub>	R <sub>o</sub>	F <sub>1<sub>o</sub></sub>
Rule-based (Sieve-based)	90.0	84.7	92.6	88.5	<b>83.8</b>	36.3	<b>94.8</b>	52.5	<b>85.3</b>	63.4	<b>97.1</b>	76.7	<b>82.9</b>	26.6	<b>98.1</b>	41.9	72.1	6.6	<b>86.7</b>	12.2
Ft-based Classifier	36.0	49.1	71.6	58.2	19.1	11.2	69.7	19.2	25.8	27.5	63.9	38.4	20.1	7.2	37.9	12.1	24.2	1.6	70.3	3.0
Ft-based BLINKout $L_{NIL}^{(F)}$	60.9	61.4	<b>95.0</b>	74.6	42.3	14.0	92.3	24.2	48.0	39.3	95.9	55.7	43.7	8.9	85.4	16.1	45.2	1.8	61.4	3.5
Th-based BM25+BERT+syn	69.5	64.9	55.8	60.0	72.3	22.8	64.5	33.7	70.2	46.2	60.7	52.5	70.6	9.7	26.2	14.1	77.4	7.0	39.2	11.9
Th-based BM25+SapBERT+syn	76.2	71.3	64.7	67.8	75.6	28.6	84.7	42.2	74.7	51.3	72.5	60.1	71.4	16.9	81.6	28.0	75.8	7.6	70.3	13.8
BLINK (BERT)	53.0	-	-	-	67.0	-	-	-	57.2	-	-	-	62.4	-	-	-	88.0	-	-	-
+ syn	58.1	-	-	-	70.4	-	-	-	60.2	-	-	-	68.8	-	-	-	87.8	-	-	-
Th-based BLINK	72.4	65.4	67.7	66.5	68.1	21.9	42.6	28.9	71.1	51.6	77.6	62.0	65.7	21.7	74.8	33.6	88.1	17.8	49.4	26.2
+ syn	80.1	74.4	76.2	75.3	72.7	25.5	54.8	34.8	69.5	47.7	60.5	53.4	71.1	22.2	48.5	30.5	87.8	16.3	44.3	23.8
NIL-rep-based BLINK	74.8	67.6	89.2	76.9	69.7	24.8	47.1	32.5	71.9	56.9	74.9	64.7	65.7	19.1	80.6	30.9	87.9	16.3	24.7	19.6
+ syn	80.3	79.8	73.5	76.5	72.6	27.8	40.7	33.0	73.8	54.4	77.1	61.7	71.4	18.6	74.8	29.8	87.9	16.2	55.7	25.0
BLINKout (- syn, BERT)	77.7	74.1	91.2	81.8	62.3	18.4	62.6	28.5	68.8	51.1	80.7	62.6	68.2	25.2	78.6	38.1	88.1	24.3	49.4	32.6
BLINKout (BERT)	86.0	85.3	87.9	86.6	70.5	33.1	58.7	42.3	75.2	55.3	77.8	64.7	69.9	33.7	84.5	48.2	87.9	21.4	72.2	33.0
BLINKout (SapBERT)	89.9	90.0	89.8	89.9	81.8	53.5	54.8	54.1	81.4	<b>72.6</b>	81.0	76.5	81.5	57.3	79.6	66.7	86.6	22.3	69.0	33.7
+ k and NIL rep tuned*	90.8	<b>91.4</b>	90.3	<b>90.9</b>	83.3	<b>58.4</b>	65.2	<b>61.6</b>	84.8	<b>72.6</b>	88.7	<b>79.8</b>	<b>82.9</b>	<b>86.2</b>	72.8	<b>79.0</b>	<b>90.7</b>	33.0	70.3	44.9
+ $L_{NIL}$	<b>91.2</b>	88.5	92.5	90.5	83.1	47.0	75.5	57.9	83.0	65.2	91.6	76.2	82.5	59.3	83.5	69.4	90.4	<b>37.1</b>	67.1	<b>47.8</b>

Table 2: Comparison results for Entity Linking with *out-of-KB* mentions (Precision, Recall,  $F_1$  scores w.r.t. *out-of-KB* mentions and overall Accuracy w.r.t. both in- and *out-of-KB* mentions). \*The parameters  $k$  and NIL rep are tuned using BLINKout (SapBERT) for ShARe/CLEF 2013 and MM datasets and BLINKout (BERT) for NILK-sample.

	ShARe/CLEF 2013			NILK-sample		
	A	F <sub>1<sub>o</sub></sub>	F <sub>1<sub>in</sub></sub>	A	F <sub>1<sub>o</sub></sub>	F <sub>1<sub>in</sub></sub>
<b>Fixed rep:</b>						
<b>NIL-rep-based BLINK+syn</b>						
NIL	70.0	52.3	75.4	87.3	3.3	88.0
NIL [ENT] It is a NIL option.	62.7	25.6	71.5	87.3	1.2	88.0
[NIL]	<b>80.3</b>	<b>76.5</b>	<b>82.0</b>	<b>87.9</b>	<b>25.0</b>	<b>90.1</b>
[NIL] [ENT] It is a NIL option.	74.5	64.9	78.0	87.8	24.7	89.4
[NIL] [ENT] It is a [NIL] option.	76.5	69.2	79.4	87.8	24.1	89.7
<b>Fine-tuned rep:</b>						
<b>BLINKout (BERT), <math>k</math> tuned</b>						
[NIL]	<b>88.7</b>	<b>89.1</b>	<b>88.5</b>	<b>90.7</b>	44.9	<b>91.8</b>
[NIL] [ENT] It is a [NIL] option.	87.6	88.2	87.2	90.3	<b>46.2</b>	91.1

Table 3: Comparison results among NIL entity representations (“rep”), either fixed or fine-tuned

**Joint Learning.** Using  $L_{NIL}$  allows to penalise the situation that a NIL mention being wrongly classified as an in-KB entity, or vice-versa. This may increase the *out-of-KB* recall  $R_o$ , but harm the precision  $P_o$  and  $F_{1o}$  scores (except for NILK-sample), as shown in Table 2. Joint learning obtained the best overall accuracy on ShARe/CLEF 2013 and best *out-of-KB*  $F_1$  for NILK-sample. More effective use of  $L_{NIL}$  warrants further studies.

**Synonyms Enhancement.** Synonyms help address the entity variant problem and further differentiate in-KB entities from NIL after the fine-tuning. Our synonym enhancement on the bi-encoder and cross-encoder greatly improved the *out-of-KB*  $F_1$  scores across most datasets over BLINK. In the zero-shot EL scenario (e.g., NILK-sample data and zero-shot test sets of other data in Appendix D), synonyms may be less useful given that synonyms of unseen entities are not presented during training. Making use of synonyms in the ZS setting warrants further studies.

## 8 RELATED WORK

**Entity Linking (EL) for KB Construction.** EL has been extensively studied, see reviews in Sevgili et al. [38], Shen et al. [39, 40]. The work in Rao et al. [34] focuses on a KB-centric view of EL,

i.e., EL as a key step in KB construction and maintenance. Most relevant to KB maintenance is the identification of *out-of-KB* or NIL mentions so that they can be placed into the KB [34]. Many recent studies in EL only considered in-KB entities [3, 4, 45]. *Out-of-KB* mention discovery is also distinct from zero-shot EL [36, 45], as the latter targets in-KB entities.

**Out-of-KB Mention and Entity Discovery.** Most studies apply the traditional threshold-based approach [6, 7, 25, 47] or NIL classification with features [34, 46, 48] and neural networks [16, 30]. We proposed NIL entity representation & classification that utilizes the embedding representation in LMs, with BERT-based EL [45].

Recent studies form NIL clusters for the resolution of *out-of-KB* mentions into potential concepts or entities [1, 20, 26]. The NIL clusters are formed either with a threshold-based approach or as a post-processing of the NIL mentions after their discovery, as summarized in Ji et al. [24]. Our methods on NIL entity representation & classification are independent of clustering methods, and we leave the combination and comparison with them for a future study.

**Out-of-KB EL Benchmarking.** We summarized most of the NIL-labeled EL datasets in the introduction, including ShARe/CLEF 2013 [41], NILK [23], NEEL 2015-2016 challenges [37], and CLEF HIPE 2020 [12]. The earliest dataset is TAC 2011 Knowledge Base Population Track [24], which has NIL mentions to enrich a KB derived from Wikipedia Infoboxes, however, is not freely available and the data size is small (around 1000 mentions in training and testing resp.). Another recent clinical note dataset is 2019-n2c2-MCN [29], which is not yet openly available for new users outside of the previous registration for the challenge. Another recent general domain dataset is EDIN-benchmark [26] to enrich older versions of Wikipedia using text mentions from news articles. EDIN-benchmark requires an adaptation set of documents, which may not be always available. We summarized and applied strategies (manual labelling, KB pruning, KB versioning) for NIL-enhanced EL data creation. As far as we know, all previous studies have considered only one of the strategies and we are the first study encompassing all three strategies for dataset construction and benchmarking.

Data w/ row ID (starts from 0)	Mention in context	Gold	$Th$ -based BLINK+syn	NIL-rep-based BLINK+syn	BLINKout (tuned)	BLINKout (tuned)+ $L_{NIL}$
ShARe/ CLEF 2013 test-2353	Discharge Diagnosis: <b>PEA arrest</b> of unclear etiology	NIL	<b>Electromechanical dis-</b> <b>sociation (C0340861)</b> <b>1.000&gt;<math>Th_{cross}</math> (0.95)</b>	<b>Electromechanical dis-</b> <b>sociation (C0340861)</b> <b>1.000</b> (where NIL 0.000, 2nd)	<b>Electromechanical dis-</b> <b>sociation (C0340861)</b> <b>0.888</b> (where NIL 0.001, 2nd)	NIL 0.718
ShARe/ CLEF 2013 test-5083	Low dose beta blockade <b>decreased HR</b> to 95 in sinus and synthroid had been restarted.	Bradycardia (C0428977)	<b>NIL (Cardiac mur-</b> <b>mur, intensity grade</b> <b>I/VI (C0232249)</b> <b>0.191&lt;<math>Th_{cross}</math> (0.95))</b> (Gold not in top-k)	<b>NIL 0.973</b> (Gold not in top-k)	Bradycardia (C0428977) 0.989	<b>NIL 0.947</b> (where Bradycardia (C0428977) 0.053, 2nd)
MM-pruned-0.1 test-14	Dietary antioxidants may play an important role in the prevention of <b>bone loss</b> and associated fractures by reducing levels of oxidative stress.	NIL	NIL (Tooth eruption disorder (C0012767) 0.627< $Th_{cross}$ (0.80))	<b>Tooth eruption disorder (C0012767) 0.426</b> (where NIL 0.320, 2nd)	NIL 1.000	NIL 0.997
MM-pruned-0.1 test-35	... including weight and fat gain, glucose intolerance, <b>hypertriglyceridemia</b> , abnormal adipocytokine levels, hypertension, and adiponectin and leptin gene expression and epigenetic changes.	NIL	<b>Hyperglycemia (C0020456)</b> <b>1.000&gt;<math>Th_{cross}</math> (0.80)</b>	<b>Hyperglycemia (C0020456)</b> <b>1.000</b> (where NIL 0.000, 3rd)	<b>Dyslipidemias (C0242339) 0.564</b> (where NIL 0.011, 3rd)	<b>Dyslipidemias (C0242339) 0.553</b> (where NIL 0.205, 3rd)
NILK-sample test-68	The ancient cities of Friesland are shown below: Dutch West Frisian Charter granted ... Workum Warkum 1399 <b>Bolsward</b> Boalsert 1455 Harlingen Harns 1234 Franeker Frjentsjer 1374	NIL	NIL (Ferdinand Bol (Q374039) 0.403< $Th_{cross}$ (0.45))	<b>Ferdinand (Q374039) 0.480</b> (where NIL 0.276, 2nd)	<b>Bol (Q2625526) 1.000</b> (where NIL 0.000, 2nd)	NIL 0.377
NILK-sample test-2054	... <b>70720 Davidskillman</b> Named in honor of David R. Skillman (b. 1945) for his decades-long contributions to asteroid searching, stellar binary star systems and as lead systems engineer for the Hubble Space Telescope at Goddard Space Flight Center. TMS Entertainment (known for "Lupin III" and "Detective Conan") would animate the popular animated series "Tiny Toon Adventures" and "Rainbow Brite". ...	NIL	<b>DAVID (Q5204342) 0.980&gt;<math>Th_{cross}</math> (0.45)</b>	<b>DAVID (Q5204342) 0.987</b> (where NIL 0.019, 2nd)	NIL 0.798	NIL 0.965
NILK-sample test-755	TMS Entertainment (known for "Lupin III" and "Detective Conan") would animate the popular animated series "Tiny Toon Adventures" and "Rainbow Brite". ...	NIL	<b>Rainbow and the Star Stealer (Q2625526) 1.000&gt;<math>Th_{cross}</math> (0.45)</b>	<b>Brite and the Star Stealer (Q2625526) 1.000</b> (where NIL 0.000, 2nd)	<b>Rainbow and the Star Stealer (Q2625526) 1.000</b> (where NIL 0.000, 2nd)	<b>Rainbow and the Star Stealer (Q2625526) 0.999</b> (where NIL 0.001, 2nd)

**Table 4: Examples of out-of-KB mention discovery from clinical notes (ShARe/CLEF 2013), biomedical publications (MedMentions), and Wikipedia articles (NILK) with the top prediction from four BERT-based Entity Linking models: Threshold-based BLINK, NIL representation based BLINK, BLINKout (tuned), and BLINKout (tuned) with joint learning ( $L_{NIL}$ ). Normalised prediction score (after softmax) and the rank (if not 1st) are displayed after the predicted entity. Wrong predictions are marked with red. For the three datasets, ShARe/CLEF 2013, MM-pruned-0.1, NILK-sample, the threshold  $Th_{cross}$  for  $Th$ -based BLINK was 0.95, 0.80, 0.45, and the tuned number of top- $k$  for BLINKout was 150, 50, 50, resp.**

Regarding evaluation, overall accuracy [7, 16, 25] is commonly used, which does not reflect the full picture. As far as we know, our study is the first to apply *out-of-KB* Precision, Recall, and  $F_1$  scores.

The research endeavour that mostly resembles ours is the ongoing work of Möller [33] on EL for KB enrichment. This study targets news events to enrich Wikipedia. Similar to the plans in Möller [33], our future work will canonicalise the new mentions by, for example, grouping and naming, and placing them in the KB.

## 9 CONCLUSION

We introduced the task of out-of-KB mention discovery from texts and proposed BLINKout, which utilizes a dynamic NIL representation & classification approach, enhanced with synonyms, founded on BERT-based Entity Linking. We also provided strategies, KB Pruning and KB Versioning, to construct out-of-KB datasets. The approach has been tested on datasets with various domains to enrich medical ontologies and WikiData. The method in this work, while extending the BERT-based Entity Linking (BLINK) approach [45], also has the potential to be applied to recent end-to-end Entity Linking methods [2]. Future studies will focus on the canonicalisation and placement of out-of-KB entities in a KB.

**Acknowledgements.** This work is supported by EPSRC projects, including ConCur (EP/V050869/1), OASIS (EP/S032347/1), UK FIRES (EP/S019111/1); and Samsung Research UK (SRUK).

## A EXPERIMENTAL DETAILS

BLINKout was implemented using PyTorch Huggingface Transformers<sup>12</sup>, based on the original BLINK library<sup>13</sup>. The BM25+BERT baseline further used the Python library rank\_bm25<sup>14</sup>. Results on the Sieve-based approach were reproduced using the original JAVA-based implementation<sup>15</sup>. All neural network models were trained using an NVIDIA Quadro RTX 8000 GPU card (48GB GPU memory).

We tuned the hyperparameters by optimising  $F_1$ , on the validation set. We generally tuned the best  $k \in \{5, 10, 20, 50, 100, 150, 200\}$  for each dataset and selected the best NIL representation. The default  $k$  for model comparison was set as 10 for Share/CLEF 2013 following Ji et al. [25] and MedMention datasets, and 4 for the NILK-sample dataset following Iurshina et al. [23]. The best BLINKout model had  $k$  as 150 for ShARe/CLEF 2013, 50 for MM-pruned-0.1, MM-2014AB, and NILK-sample, and 100 for MM-pruned-0.2. The NIL representation was [NIL] for the datasets except [NIL] with [NIL]-definition for NILK-sample (only when using the fixed setting) and MM-pruned-0.1. For the threshold-based BLINK method, the threshold  $Th_{cross}$  was 0.95, 0.55, 0.80, 0.95, 0.45 for the datasets ShARe/CLEF 2013, MM-2014AB, MM-pruned-0.1, MM-pruned-0.2, and NILK-sample, resp.

<sup>12</sup><https://huggingface.co/docs/transformers/index>

<sup>13</sup><https://github.com/facebookresearch/BLINK>


<sup>14</sup><https://pypi.org/project/rank-bm25/>

<sup>15</sup><https://github.com/jennydsuza9/disorder-normalizer>



	ShARe/CLEF 2013			MM-pruned-0.1			MM-pruned-0.2			MM-2014AB			NILK-sample		
	$P_{in}$	$R_{in}$	$F_{1_{in}}$	$P_{in}$	$R_{in}$	$F_{1_{in}}$	$P_{in}$	$R_{in}$	$F_{1_{in}}$	$P_{in}$	$R_{in}$	$F_{1_{in}}$	$P_{in}$	$R_{in}$	$F_{1_{in}}$
Rule-based (Sieve-based)	<b>93.1</b>	88.8	90.9	<b>95.9</b>	82.9	<b>88.9</b>	<b>95.5</b>	82.2	<b>88.4</b>	<b>96.1</b>	82.1	<b>88.6</b>	88.2	71.9	79.2
Ft-based Classifier	19.0	24.4	21.4	26.6	14.9	19.0	24.1	15.8	19.1	25.0	19.2	21.7	70.9	23.5	35.3
Ft-based BLINKout $L_{NIL}^{(Ft)}$	60.4	44.6	51.3	72.1	38.1	49.9	57.0	35.5	43.7	78.1	41.4	54.1	88.6	45.0	59.7
Th-based BM25+BERT+syn	71.3	76.3	73.6	86.2	73.0	79.0	79.2	72.6	75.8	80.5	73.0	76.6	83.8	78.0	80.8
Th-based BM25+SapBERT+syn	78.2	81.7	79.9	88.7	75.1	81.3	84.4	75.3	79.6	89.7	70.9	79.2	86.6	75.9	80.9
BLINK (BERT)	53.0	78.4	63.2	67.0	72.6	69.7	57.2	72.2	63.8	62.4	65.8	64.1	88.0	89.3	88.6
+ syn	58.1	85.8	69.3	70.4	76.3	73.2	60.2	76.0	67.2	68.8	72.6	70.6	87.8	89.1	88.4
Th-based BLINK	76.0	74.7	75.3	76.3	70.2	73.2	79.9	69.3	74.2	75.3	65.3	69.9	91.1	88.7	89.9
+ syn	82.9	81.9	82.4	82.2	74.2	78.0	77.2	71.8	74.4	77.3	72.3	74.7	90.8	88.5	89.6
NIL-rep-based BLINK	80.2	68.0	73.6	77.4	71.6	74.4	77.5	71.0	74.1	78.8	64.9	71.2	87.9	89.6	88.9
+ syn	80.5	83.6	82.0	78.3	75.3	76.8	81.0	74.5	77.6	85.3	71.3	77.7	91.8	88.4	90.1
BLINKout (-syn, BERT)	80.2	71.3	75.4	77.9	62.2	69.2	77.4	65.6	71.0	76.5	67.6	71.8	90.1	88.7	89.4
BLINKout (BERT)	86.4	85.2	85.8	76.5	71.5	73.9	83.4	74.5	78.7	75.3	69.1	72.1	91.5	88.2	89.8
BLINKout (SapBERT)*	89.9	90.0	90.0	84.3	<b>84.1</b>	84.2	84.1	81.5	82.8	83.4	81.6	82.5	89.7	86.9	88.3
+ k and NIL rep tuned	90.5	<b>91.0</b>	90.7	87.1	83.3	85.1	88.9	<b>83.7</b>	86.2	82.8	<b>83.5</b>	83.1	<b>92.6</b>	<b>91.0</b>	<b>91.8</b>
+ $L_{NIL}$	92.5	90.5	<b>91.5</b>	88.2	83.7	85.9	90.3	80.7	85.2	84.3	82.4	83.4	91.9	90.8	91.3

**Table 5: Comparison results for in-KB Entity Linking (Precision, Recall,  $F_1$  scores for in-KB entities). \* The parameters  $k$  and NIL rep are tuned using BLINKout (SapBERT) for ShARe/CLEF 2013 and MM datasets, and BLINKout (BERT) for NILK-sample.**



**Figure 3: Out-of-KB  $F_1$ , in-KB  $F_1$ , and overall accuracy w.r.t. the number of top- $k$  candidates using BLINKout (SapBERT) on MM-2014AB validation set**

We followed Wu et al. [45] to use a maximum of 32 tokens for a mention in context and 128 for an entity candidate (including its synonyms, 32 for NILK-sample), learning rate as  $3e-05$ , bi-encoder training batch size as 16, cross-encoder training batch size as 1. We used the “large” version of LM (e.g., BERT-large, 680M parameters) in the bi-encoder and “base” version of LM (e.g., BERT-base, 110M parameters) in the cross-encoder. For SapBERT and PubMedBERT, we used “base” version for the bi-encoder as no “large” is available.

The parameter in joint learning,  $\lambda_{NIL}$ , was tuned to 0.25 (as in [16]) for ShARe/CLEF 2013 and MM-pruned-0.2 datasets, to 0.2 for MM-pruned-0.1, 0.05 for MM-2014AB, and 0.01 for NILK-sample, based on the validation set. We find that value of  $\lambda_{NIL}$  is near to the percentage of out-of-KB mentions in Table 1 for each dataset.

We optimised the bi-encoder and cross-encoder using AdamW [28], with 3 and 4 epochs, resp. (except for NILK-sample, 1 epoch each due to the large data size). We trained all models with fixed

$k=10$	ShARe/CLEF 2013-ZS			MM-pruned-0.1-ZS		
	A	$F_{1_o}$	$F_{1_{in}}$	A	$F_{1_o}$	$F_{1_{in}}$
BLINKout (-syn, BERT)	78.0	<b>90.7</b>	41.3	<b>48.2</b>	45.1	<b>49.7</b>
BLINKout (BERT)	<b>80.3</b>	90.2	<b>54.9</b>	43.8	<b>51.6</b>	41.0
Diff in ZS setting	+2.3	-0.5	+13.6	-4.4	+6.5	-8.7
Diff in original setting	+7.3	+4.8	+10.4	+8.2	+13.8	+4.7

**Table 6: Comparison results on zero-shot (ZS) testing sets on BLINKout w.r.t. synonym enhancement. Their differences (“Diff”) in the original data setting is based on Table 2 and 5.**

random seeds to obtain reproducible results and the actual variance was low with random seeds, less than 1% (e.g., 0.85% on MM-2014AB with BLINKout,  $k=10$ , over three runs).

**Time Used in Training and Inference.** The training of bi-encoder with SapBERT on ShARe/CLEF 2013 took approx. 0.63 hour; cross-encoder with SapBERT took approx. 0.88 hour when  $k=10$  and approx. 8.67 hours when  $k=150$ . The inferencing stage is very fast compared to the training: inferencing with a new mention took approx. 0.06 and 0.57 second for  $k=10$  and  $k=150$ , resp.

## B RESULTS W.R.T. NUMBER OF TOP- $k$

Figure 3 displays the change of  $F_{1_o}$ ,  $F_{1_{in}}$ , and  $A$  with respect to the number of  $k$  on MM-2014AB with the BLINKout model (SapBERT).

## C RESULTS FOR IN-KB ENTITIES

In-KB EL results on the four datasets are presented in Table 5.

## D ZERO-SHOT EL AND SYNONYMS

Zero-shot entity linking results w.r.t. the use of synonyms to enhance BLINKout are presented in Table 6. Results show that in the zero-shot setting of the testing sets, i.e., after removing testing mentions having overlapped entities to the training set, the performance of out-of-KB mention discovery (e.g.,  $F_{1_o}$ ) with BLINKout is less sensitive to synonym enhancement. The inference of out-of-KB entities may be affected when ranking the unseen entities, which have abundant synonyms not presented during training.

## REFERENCES

- [1] Dhruv Agarwal, Rico Angell, Nicholas Monath, and Andrew McCallum. 2021. Entity Linking and Discovery via Arborescence-based Supervised Clustering. *CoRR abs/2109.01242* (2021), 1–12. arXiv:2109.01242 <https://arxiv.org/abs/2109.01242>
- [2] Tom Ayoola, Shubhi Tyagi, Joseph Fisher, Christos Christodoulopoulos, and Andrea Pierleoni. 2022. ReFinED: An Efficient Zero-shot-capable Approach to End-to-End Entity Linking. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*. Association for Computational Linguistics, Hybrid: Seattle, Washington + Online, 209–220. <https://doi.org/10.18653/v1/2022.naacl-industry.24>
- [3] 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)*. Association for Computational Linguistics, Online, 3122–3137. <https://doi.org/10.18653/v1/2020.emnlp-main.253>
- [4] Rajarshi Bhowmik, Karl Stratos, and Gerard de Melo. 2021. Fast and Effective Biomedical Entity Linking Using a Dual Encoder. In *Proceedings of the 12th International Workshop on Health Text Mining and Information Analysis*. Association for Computational Linguistics, online, 28–37. <https://aclanthology.org/2021.louhi-1.4>
- [5] Olivier Bodenreider. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research* 32, suppl\_1 (01 2004), D267–D270. <https://doi.org/10.1093/nar/gkh061>
- [6] Razvan Bunescu and Marius Pasca. 2006. Using Encyclopedic Knowledge for Named entity Disambiguation. In *11th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Trento, Italy, 9–16. <https://aclanthology.org/E06-1002>
- [7] Lihu Chen, Gaël Varoquaux, and Fabian M. Suchanek. 2021. A Lightweight Neural Model for Biomedical Entity Linking. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 14 (May 2021), 12657–12665. <https://doi.org/10.1609/aaai.v35i14.17499>
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [9] Kevin Donnelly et al. 2006. SNOMED-CT: The advanced terminology and coding system for eHealth. In *Medical and Care Computatics 3. Studies in health technology and informatics*, Vol. 121. IOS Press, Amsterdam, Netherlands, 279–290.
- [10] Mark Dredze, Paul McNamee, Delip Rao, Adam Gerber, and Tim Finin. 2010. Entity Disambiguation for Knowledge Base Population. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*. Coling 2010 Organizing Committee, Beijing, China, 277–285. <https://aclanthology.org/C10-1032>
- [11] Jennifer D'Souza and Vincent Ng. 2015. Sieve-Based Entity Linking for the Biomedical Domain. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Beijing, China, 297–302. <https://doi.org/10.3115/v1/P15-2049>
- [12] Maud Ehrmann, Matteo Romanello, Alex Flückiger, and Simon Clematide. 2020. Extended overview of CLEF HIPE 2020: named entity processing on historical newspapers. In *Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum*. CEUR Workshop Proceedings (CEUR-WS.org), Thessaloniki, Greece, 1–38.
- [13] Noémie Elhadad, Wendy Chapman, and Guergana Savova. 2013. ShARE-CLEF eHealth 2013: Natural Language Processing and Information Retrieval for Clinical Care 1.0. <https://physionet.org/content/shareclefehealth2013/1.0/Task1ShAREGuidelines2013.pdf>.
- [14] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6894–6910. <https://doi.org/10.18653/v1/2021.emnlp-main.552>
- [15] Thomas R. Gruber. 1995. Toward principles for the design of ontologies used for knowledge sharing? *International Journal of Human-Computer Studies* 43, 5 (1995), 907–928. <https://doi.org/10.1006/ijhc.1995.1081>
- [16] Yingjie Gu, Xiaoye Qu, Zhefeng Wang, Baoxing Huai, Nicholas Jing Yuan, and Xiaolin Gui. 2021. Read, Retrospect, Select: An MRC Framework to Short Text Entity Linking. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 14 (May 2021), 12920–12928. <https://doi.org/10.1609/aaai.v35i14.17528>
- [17] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Trans. Comput. Healthcare* 3, 1, Article 2 (oct 2021), 23 pages. <https://doi.org/10.1145/3458754>
- [18] Yuan He, Jiaoyan Chen, Hang Dong, Ian Horrocks, Carlo Allocca, Taehun Kim, and Brahmananda Sapkota. 2023. DeepOnto: A Python Package for Ontology Engineering with Deep Learning. *arXiv preprint arXiv:2307.03067* (2023).
- [19] Yuan He, Jiaoyan Chen, Hang Dong, Ernesto Jiménez-Ruiz, Ali Hadian, and Ian Horrocks. 2022. Machine learning-friendly biomedical datasets for equivalence and subsumption ontology matching. In *International Semantic Web Conference*. Springer, Cham, Switzerland, 575–591.
- [20] Nicolas Heist and Heiko Paulheim. 2023. NASTyLinker: NIL-Aware Scalable Transformer-Based Entity Linker. In *The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings (Lecture Notes in Computer Science, Vol. 13870)*, Ernesto Jiménez-Ruiz, Jamie P. McCusker, Daniel Faria, Mauro Dragoni, Anastasia Dimou, Raphaël Troncy, and Sven Hertling (Eds.). Springer, Cham, 174–191.
- [21] Johannes Hoffart, Yasemin Altun, and Gerhard Weikum. 2014. Discovering Emerging Entities with Ambiguous Names. In *Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14)*. Association for Computing Machinery, New York, NY, USA, 385–396. <https://doi.org/10.1145/2566486.2568003>
- [22] Johannes Hoffart, Dragan Milchevski, Gerhard Weikum, Avishek Anand, and Jaspreet Singh. 2016. The Knowledge Awakens: Keeping Knowledge Bases Fresh with Emerging Entities. In *Proceedings of the 25th International Conference Companion on World Wide Web (Montréal, Québec, Canada) (WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 203–206. <https://doi.org/10.1145/2872518.2890537>
- [23] Anastasia Iurshina, Jiaxin Pan, Rafika Boutalbi, and Steffen Staab. 2022. NILK: Entity Linking Dataset Targeting NIL-Linking Cases. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 4069–4073. <https://doi.org/10.1145/3511808.3557659>
- [24] Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. 2011. Overview of the TAC 2011 knowledge base population track. In *Third text analysis conference (TAC 2011)*. National Institute of Standards and Technology, Gaithersburg, Maryland, USA, 1–33.
- [25] Zongcheng Ji, Qiang Wei, and Hua Xu. 2020. BERT-based ranking for biomedical entity normalization. *AMIA Summits on Translational Science Proceedings* 2020 (2020), 269.
- [26] Nora Kassner, Fabio Petroni, Mikhail Plekhanov, Sebastian Riedel, and Nicola Cancedda. 2022. EDIN: An End-to-end Benchmark and Pipeline for Unknown Entity Discovery and Indexing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 8659–8673.
- [27] Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella, and Nigel Collier. 2021. Self-Alignment Pretraining for Biomedical Entity Representations. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 4228–4238. <https://doi.org/10.18653/v1/2021.naacl-main.334>
- [28] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, OpenReview.net, New Orleans, LA, USA, 1–10*.
- [29] Yen-Fu Luo, Sam Henry, Yanshan Wang, Feichen Shen, Ozlem Uzuner, and Anna Rumshisky. 2020. The 2019 n2c2/UMass Lowell shared task on clinical concept normalization. *Journal of the American Medical Informatics Association* 27, 10 (09 2020), 1529–e1. <https://doi.org/10.1093/jamia/ocaa106>
- [30] Pedro Henrique Martins, Zita Marinho, and André F. T. Martins. 2019. Joint Learning of Named Entity Recognition and Entity Linking. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. Association for Computational Linguistics, Florence, Italy, 190–196. <https://doi.org/10.18653/v1/P19-2026>
- [31] Paul McNamee, Mark Dredze, Adam Gerber, Nikesha Garera, Tim Finin, James Mayfield, Christine Piatko, Delip Rao, David Yarowsky, Markus Dreyer, et al. 2009. HLTCOE Approaches to Knowledge Base Population at TAC 2009. In *Proceedings of the 2009 Text Analysis Conference*. National Institute of Standards and Technology, Gaithersburg, Maryland, USA, 1–10.
- [32] Sunil Mohan and Donghui Li. 2019. MedMentions: A Large Biomedical Corpus Annotated with {UMLS} Concepts. In *Automated Knowledge Base Construction (AKBC)*. openreview.net, Amherst, MA, USA, 1–13. <https://doi.org/10.24432/C5G59C>
- [33] Cedric Möller. 2022. Knowledge Graph Population with Out-of-KG Entities. In *The Semantic Web: ESWC 2022 Satellite Events - Hersonissos (Lecture Notes in Computer Science, Vol. 13384)*, Paul Groth, Anisa Rula, Jodi Schneider, Ilaria Tiddi, Elena Simperl, Panos Alexopoulos, Rinke Hoekstra, Mehwish Alam, Anastasia Dimou, and Minna Tamper (Eds.). Springer, Hersonissos, Crete, Greece, 199–214. [https://doi.org/10.1007/978-3-031-11609-4\\_35](https://doi.org/10.1007/978-3-031-11609-4_35)
- [34] Delip Rao, Paul McNamee, and Mark Dredze. 2013. Entity Linking: Finding Extracted Entities in a Knowledge Base. In *Multi-source, Multilingual Information Extraction and Summarization*, Thierry Poibeau, Horacio Saggion, Jakub Piskorski,

- and Roman Yangarber (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 93–115. [https://doi.org/10.1007/978-3-642-28569-1\\_5](https://doi.org/10.1007/978-3-642-28569-1_5) Series Title: Theory and Applications of Natural Language Processing.
- [35] Nils Reimers and Iryna Gurevych. 2019. 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)*. Association for Computational Linguistics, Hong Kong, China, 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
- [36] Petar Ristoski, Zhizhong Lin, and Qunzhi Zhou. 2021. KG-ZESHEL: Knowledge Graph-Enhanced Zero-Shot Entity Linking. In *Proceedings of the 11th on Knowledge Capture Conference (Virtual Event, USA) (K-CAP '21)*. Association for Computing Machinery, New York, NY, USA, 49–56. <https://doi.org/10.1145/3460210.3493549>
- [37] Giuseppe Rizzo, Bianca Pereira, Andrea Varga, Marieke Van Erp, and Amparo Elizabeth Cano Basave. 2017. Lessons learnt from the Named Entity rEcognition and Linking (NEEL) challenge series. *Semantic Web* 8, 5 (2017), 667–700.
- [38] Özge Sevgili, Artem Shelmanov, Mikhail Arkhipov, Alexander Panchenko, and Chris Biemann. 2022. Neural entity linking: A survey of models based on deep learning. *Semantic Web* 13, 3 (Jan. 2022), 527–570. <https://doi.org/10.3233/SW-222986> Publisher: IOS Press.
- [39] Wei Shen, Yuhan Li, Yanan Liu, Jiawei Han, Jianyong Wang, and Xiaojie Yuan. 2023. Entity Linking Meets Deep Learning: Techniques and Solutions. *IEEE Transactions on Knowledge and Data Engineering* 35, 3 (2023), 2556–2578. <https://doi.org/10.1109/TKDE.2021.3117715>
- [40] Wei Shen, Jianyong Wang, and Jiawei Han. 2014. Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* 27, 2 (2014), 443–460.
- [41] Hanna Suominen, Sanna Salanterä, Sumithra Velupillai, Wendy Webber Chapman, Guergana K. Savova, Noemie Elhadad, Sameer Pradhan, Brett R. South, Danielle L. Mowery, Gareth J. F. Jones, Johannes Leveling, Liadh Kelly, Lorraine Goeuriot, David Martínez, and Guido Zuccon. 2013. Overview of the ShARe/CLEF eHealth Evaluation Lab 2013. In *Information Access Evaluation. Multilinguality, Multimodality, and Visualization - 4th International Conference of the CLEF Initiative, CLEF 2013 (Lecture Notes in Computer Science, Vol. 8138)*, Pamela Forner, Henning Müller, Roberto Paredes, Paolo Rosso, and Benno Stein (Eds.). Springer, Valencia, Spain, 212–231. [https://doi.org/10.1007/978-3-642-40802-1\\_24](https://doi.org/10.1007/978-3-642-40802-1_24)
- [42] Stephen RF Twigg, Robert B Hufnagel, Kerry A Miller, Yan Zhou, Simon J McGowan, John Taylor, Jude Craft, Jenny C Taylor, Stephanie L Santoro, Taosheng Huang, et al. 2016. A recurrent mosaic mutation in SMO, encoding the hedgehog signal transducer smoothed, is the major cause of Curry-Jones syndrome. *The American Journal of Human Genetics* 98, 6 (2016), 1256–1265.
- [43] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM* 57, 10 (sep 2014), 78–85. <https://doi.org/10.1145/2629489>
- [44] World Health Organization. 2022. Tracking SARS-CoV-2 variants. <https://www.who.int/activities/tracking-SARS-CoV-2-variants>.
- [45] Ledell Wu, Fabio Petroni, Martin Josifovski, Sebastian Riedel, and Luke Zettlemoyer. 2020. Scalable Zero-shot Entity Linking with Dense Entity Retrieval. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 6397–6407. <https://doi.org/10.18653/v1/2020.emnlp-main.519>
- [46] Zhaohui Wu, Yang Song, and C. Lee Giles. 2016. Exploring Multiple Feature Spaces for Novel Entity Discovery. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, Dale Schuurmans and Michael P. Wellman (Eds.). AAAI Press, Phoenix, Arizona, USA, 3073–3079. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12261>
- [47] Dongfang Xu and Timothy Miller. 2022. A simple neural vector space model for medical concept normalization using concept embeddings. *Journal of Biomedical Informatics* 130 (2022), 104080.
- [48] Shuo Zhang, Edgar Meij, Krisztian Balog, and Ridho Reinanda. 2020. Novel Entity Discovery from Web Tables. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 1298–1308. <https://doi.org/10.1145/3366423.3380205>