

Proxy-based Zero-Shot Entity Linking by Effective Candidate Retrieval

Maciej Wiatrak^{1*}, Eirini Arvaniti¹, Angus Brayne¹, Jonas Vetterle^{1,2}, Aaron Sim¹

¹BenevolentAI ²Moonfire Ventures

London, United Kingdom

{maciej.wiatrak, eirini.arvaniti, angus.brayne, aaron.sim}@benevolent.ai
jonas@moonfire.com

Abstract

A recent advancement in the domain of biomedical Entity Linking is the development of powerful two-stage algorithms – an initial *candidate retrieval* stage that generates a shortlist of entities for each mention, followed by a *candidate ranking* stage. However, the effectiveness of both stages are inextricably dependent on computationally expensive components. Specifically, in candidate retrieval via dense representation retrieval it is important to have *hard* negative samples, which require repeated forward passes and nearest neighbour searches across the entire entity label set throughout training. In this work, we show that pairing a proxy-based metric learning loss with an adversarial regularizer provides an efficient alternative to hard negative sampling in the candidate retrieval stage. In particular, we show competitive performance on the recall@1 metric, thereby providing the option to leave out the expensive candidate ranking step. Finally, we demonstrate how the model can be used in a zero-shot setting to discover out of knowledge base biomedical entities.

1 Introduction

The defining challenge in biomedical Entity Linking (EL) is performing classification over a large number of entity labels with limited availability of labelled mention data, in a constantly evolving knowledge base. For instance, while the Unified Medical Language System (UMLS) knowledge base (Bodenreider, 2004) contains millions of unique entity labels, the EL training data in the biomedical domain as a whole is notoriously scarce, particularly when compared to the general domain – Wikipedia, for instance, is powerful as *both* a Knowledge base and a source of matching entities and mentions. Furthermore, biomedical knowledge bases are evolving rapidly with new entities being added constantly. Given this knowledge base

evolution and scarcity of training data it is crucial that biomedical entity linking systems can scale efficiently to large entity sets, and can discover or discern entities outside of the knowledge base and training data.

Recent methods in the general entity linking domain (Logeswaran et al., 2019; Wu et al., 2020) address the data issue with zero-shot entity linking systems that use entity descriptions to form entity representations and generalise to entities without mentions. A particularly powerful architecture was initially proposed by Humeau et al. (2019) and further improved by Wu et al. (2020). It consists of a two-stage approach: 1) candidate retrieval in a dense space performed by a *bi-encoder* (Wu et al., 2020) which independently embeds the entity mention and its description, and 2) candidate ranking performed by a *cross-encoder* which attends across both the mention and entity description (Logeswaran et al., 2019). In this work we focus on the former, which is traditionally optimised with the cross-entropy (CE) loss and aims to maximise the similarity between the entity mention and its description relative to the similarities of incorrect mention-description pairs. In practice, the large number of knowledge base entities necessitates the use of negative sampling to avoid the computational burden of comparing each mention to all of the entity descriptions. However, if the sampled distribution of negatives is not reflective of the model distribution, the performance may be poor. Recently, Zhang and Stratos (2021) showed that using hard negatives - the highest scoring incorrect examples - results in bias reduction through better approximation of the model distribution. Collecting hard negatives is computationally expensive, as it requires periodically performing inference and retrieving approximate nearest neighbours for each mention.

At the ranking stage, negative sampling is not required, as the number of candidates usually does

* Corresponding author.

not exceed 64. However, the state-of-the-art cross-encoder model used for ranking is very expensive to run, scaling quadratically with the input sequence length. This highlights the need for efficient and performant candidate retrieval models capable of disambiguating mentions without the need for the expensive ranking step.

In this paper, we propose and evaluate a novel loss for the candidate retrieval model, which breaks the dependency between the positive and negative pairs. Our contributions are: (1) a novel loss which significantly outperforms the benchmark cross-entropy loss on the candidate retrieval task when using random negatives, and performs competitively when using hard negatives. (2) We design and apply an adversarial regularization method, based on the Fast Gradient Sign Method (Goodfellow et al., 2015), which is designed to simulate hard negative samples without expensively mining them. (3) We construct a biomedical dataset for out of knowledge base detection evaluation using the MedMentions corpus and show that our model can robustly identify mentions that lack a corresponding entry in the knowledge base, while maintaining high performance on the retrieval task.

Our main testing ground is the biomedical entity linking dataset MedMentions (Mohan and Li, 2019), which utilizes UMLS as its knowledge base. Additionally, to confirm that our method works also in the general, non-biomedical domain, we evaluate it on the Zero-Shot Entity Linking (ZESHEL) dataset proposed in Logeswaran et al. (2019). We focus on the retrieval task with the recall@1 metric, because we are aiming to predict the entity directly without requiring the additional expensive ranking stage. Our results show that both the proposed loss and regularization improve performance, achieving state-of-the-art results on recall@1 and competitive performance on recall@64 on both datasets. Finally, we demonstrate that our model can robustly identify biomedical out of knowledge base entities, without requiring any changes to the training procedure.

2 Related Work

Zero-Shot Entity Linking There is a plethora of work on zero-shot entity linking methods leveraging the bi-encoder architecture (Wu et al., 2020) for candidate retrieval. These include novel scoring functions between the input and the label (Humeau et al., 2019; Luan et al., 2021; Khattab and Zaharia,

2020), cross-domain pretraining methods (Varma et al., 2021), training and inference optimisation techniques (Bhowmik et al., 2021) and effective entity representation methods (Ma et al., 2021). Our work instead focuses on optimising the candidate retriever’s loss function.

The impact of hard negatives on the entity linking model performance has also been investigated (Gillick et al., 2019; Zhang and Stratos, 2021). Notably, Zhang and Stratos (2021) develop analytical tools to explain the role of hard negatives and evaluate their model on the zero-shot entity linking task. We draw on this work, but move away from the CE loss towards a novel contrastive proxy-based loss.

Finally, there is a body of work on zero-shot entity linking in the biomedical domain using clustering (Angell et al., 2021; Agarwal et al., 2021). Our method does not consider the affinities between mentions directly and links them independently. Therefore, we do not study entity discovery.

An important aspect of biomedical entity linking systems is the detection of “unlinkable” mentions that lack a corresponding entry in the Knowledge Base - referred to as NIL detection. Methods for this task can be grouped into four main strategies (Shen et al., 2014; Sevgili et al., 2020): (1) label a mention as NIL when the corresponding candidate retriever does not return any candidate entities (Tsai and Roth, 2016), (2) assign the NIL label to mentions whose corresponding top-ranked entity does not exceed some score threshold (Bunescu and Pasca, 2006; Gottipati and Jiang, 2011; Lazic et al., 2015), (3) train a classifier that predicts whether the top-ranked entity for a given mention is correct (Moreno et al., 2017), (4) explicitly introduce a NIL class to the candidate ranking model (Kolitsas et al., 2018). A downside of the final approach is that knowledge of the NIL mention distribution is required at training time. In this work we tune a NIL score threshold (2) on a validation set. Detecting unlinkable mentions is particularly important in the biomedical domain, where the knowledge bases are rapidly evolving.

Proxy-based Losses State-of-the-art entity linking models such as BLINK (Wu et al., 2020) leverage metric learning loss during training to make mentions similar to its assigned entity representations. Metric learning losses could be divided into two categories, pair-based and proxy-based losses (Kim et al., 2020). Pair-based losses can leverage semantic relations between data points, here

mentions. However, training them can be highly computationally expensive. On the other hand, proxy-based losses are significantly less computationally complex. This is done by establishing a proxy for each class and trying to increase the similarity between data points and its assigned proxies. Therefore, avoiding comparing the mentions to each other in favour of comparing the mentions to their proxies. We draw heavily on proxy-based losses (Movshovitz-Attias et al., 2017; Kim et al., 2020) from metric learning by treating entity descriptions as the proxies. We establish a proxy for each entity, creating mention-proxy (i.e. entity) pairs, and optimise the model to embed the mention close to its assigned proxy. The loss proposed here is similar to the Proxy-NCA loss of Movshovitz-Attias et al. (2017). Our modification is the use of the Softplus function, similar to Kim et al. (2020), to avoid a vanishing gradient for the true mention-proxy pair.

Adversarial Regularization Entity linking systems often rely on careful mining of hard negative examples to boost their performance (Gillick et al., 2019; Zhang and Stratos, 2021) at the expense of increased computational complexity. The model needs update hard negatives for each mention periodically. A potential alternative to hard negative mining is training on adversarial examples (Szegedy et al., 2013; Goodfellow et al., 2015) - synthetic data points designed to induce the model to making incorrect predictions, such that they are more challenging. Adversarial training can be seen as data augmentation and can help reduce overfitting. Goodfellow et al. (2015) introduced a simple method for generating adversarial examples, called Fast Gradient Sign Method (FGSM), which we build upon in this work. FGSM creates adversarial examples by applying small perturbations to the original inputs - often the word embeddings for NLP problems. FGSM has been used successfully as a regulariser in supervised and semi-supervised NLP tasks (Miyato et al., 2016; Pan et al., 2021). Here, we follow a similar approach and use FGSM to augment our training pairs with adversarial positive and negative examples.

3 Task formulation

In the **Entity Linking task** we are provided with a list of documents $D \in \mathcal{D}$, where each document has a set of mentions $M_D = \{m_1, m_2, \dots, m_{N_D}\}$. The task is to link each mention m_i to an entity

e_i , where each entity belongs to the Knowledge Base (KB) \mathcal{E} . In this work we focus specifically on the problem of biomedical zero-shot entity linking. The setup for the zero-shot task is the same as for entity linking introduced above, except that the set of entities present in the test set is not present in the training set, i.e. $\mathcal{E}_{\text{train}} \cap \mathcal{E}_{\text{test}} = \emptyset$ with $\mathcal{E}_{\text{train}} \cup \mathcal{E}_{\text{test}} = \mathcal{E}$. We focus specifically on the **Candidate Retrieval** task, where the goal is given a mention m_i , reduce the pool of potential candidate entities from a KB to a smaller subset. Candidate retrieval is crucial for biomedical entity linking because of the large size of knowledge bases. In this work we use the bi-encoder architecture for candidate retrieval. Finally, in addition to the *in-KB* entity linking task, where you only consider entities inside the KB, we also consider an **out of KB** scenario, where the task is to map mentions to the augmented set of labels $\mathcal{E} \cup \text{NIL}$, with NIL indicating the absence of a corresponding KB entity.

4 Methods

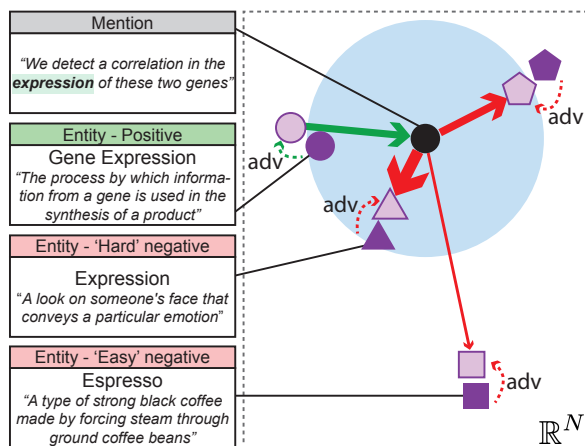


Figure 1: Overview of our proxy-based entity linking method. The mention and entity embeddings are encoded into a joint embedding space. During training, the magnitude of the gradients of the Proxy loss function with respect to the embedding coordinates is a function of the similarity between the mention and the entities (proxies). The gradients are represented by arrows whose widths indicate their magnitude. The adv-labelled dotted arrows are the Fast Gradient Sign Method adversarial perturbations. The blue circle symbolizes the margin δ .

In this section, we review the categorical CE loss, used by current state-of-the-art models, in the context of entity linking (Wu et al., 2020; Zhang and Stratos, 2021). We then compare it to our proposed Proxy-based loss. Finally, we describe

and motivate our regularization approach.

4.1 Loss

Given a set of data points corresponding to mention representations $m \in M$ and to a set of proxies corresponding to entities $e \in \mathcal{E}$, the categorical CE loss is defined as:

$$L_{\text{CE}}(m, P) := -\log \left(\frac{\exp(s(m, e^+))}{\sum_{e \in P} \exp(s(m, e))} \right), \quad (1)$$

where $s(\cdot, \cdot)$ denotes a similarity function (e.g. cosine similarity or dot product), e^+ is the positive proxy for mention representation m , P^- is a set of negative proxies used as negative samples, and $P = \{e^+\} \cup P^-$.

The gradient of the CE loss with respect to $s(m, e)$ is given by:

$$\frac{\partial L_{\text{CE}}}{\partial s(m, e)} = \begin{cases} -1 + \frac{\exp(s(m, e^+))}{\sum_{e \in P} \exp(s(m, e))}, & e = e^+ \\ \frac{\exp(s(m, e^-))}{\sum_{e \in P} \exp(s(m, e))}, & e \in P^- \end{cases} \quad (2)$$

In practice training is performed with negative sampling. If the negatives are sampled randomly, often the exponential term for the positive entity is much larger than that of the negative samples and the gradients vanish. When $s(m, e^+) \gg s(m, e^-) \forall e^- \in P^-$ then $\partial L_{\text{CE}} / \partial s(m, e) \rightarrow 0$. This behaviour is desirable when training with the full distribution of negative pairs, but stifles learning in the noisier sampling setup. A common approach is the use of hard negatives (Gillick et al., 2019; Zhang and Stratos, 2021), which increases performance over training with random negatives at the cost of increased computational complexity.

On the other hand, contrastive metric learning losses (Bromley et al., 1993; Chopra et al., 2005; Hadsell et al., 2006) alleviate the vanishing gradients problem by decoupling the positive and negative loss terms. Proxy-based contrastive losses, such as Proxy-NCA (Movshovitz-Attias et al., 2017), aim to increase the similarity between a data point x and its assigned proxy e^+ , while decreasing the similarity between x and its negative proxies $e^- \in P^-$. As demonstrated in (Kim et al., 2020), a downside of Proxy-NCA is that the scale of its gradient is constant for positive samples. This issue is alleviated by the Proxy Anchor loss (Kim et al., 2020), whose gradient reflects the rel-

ative hardness of both positive and negative pairs, resulting in improved model performance.

Drawing inspiration from the proxy-based metric learning losses described above, we formulate our Proxy-based (Pb) candidate retrieval loss as follows:

$$L_{\text{Pb}}(m, P) = \log(1 + \exp(-\alpha(s(m, e^+) - \delta))) + \log(1 + \sum_{e^- \in P^-} \exp(\alpha(s(m, e^-) + \delta))), \quad (3)$$

where we use the same notation as in Eq. 1. In addition, α is a hyperparameter controlling how strongly positive and negative samples pull and push each other, and δ is a margin. If α and δ are large, the model will be strongly penalized for the positive pair being too far from each other, and conversely the negative pair for being too close to each other. If α and δ are small, the model will receive weaker feedback. The Softplus function, a smooth approximation of the ReLU, introduces an additional margin beyond which the model stops penalising both positive and negative pairs, thus reducing overfitting. The gradient of our Proxy-based loss function is given by:

$$\frac{\partial L_{\text{Pb}}}{\partial s(m, e)} = \begin{cases} \frac{-\alpha \exp(-\alpha s^+)}{1 + \exp(-\alpha s^+)}, & e = e^+ \\ \frac{\alpha \exp(\alpha s^-)}{1 + \sum_{e^- \in P^-} \exp(\alpha s^-)}, & e \in P^- \end{cases} \quad (4)$$

where $s^+ = s(m, e^+) - \delta$, $s^- = s(m, e^-) + \delta$. This gradient reflects the relative hardness of negative examples, decoupled from the positive pair, which makes it less sensitive to the choice of negative sampling scheme.

4.2 Regularization

Our regularization approach is based on a simple adversarial training technique, called *Fast Gradient Sign Method* (FGSM) (Goodfellow et al., 2015). The idea of FGSM is to generate adversarial examples according to the following equation:

$$x_{\text{adv}} = x + \epsilon * \text{sign}(\nabla_x L(x, y)) \quad (5)$$

where x is the original training example, y its corresponding label, L the loss function that is minimised during model training, and ϵ a small number defining the magnitude of the perturbation.

FGSM applies a small perturbation to the input example that should not change the label of the resulting example x_{adv} . However, Goodfellow et al. (2015) demonstrated that even infinitesimal perturbations can cause drastic changes to the model output when carefully designed. This effect is due to the locally linear nature of neural networks in combination with the high dimensionality of their inputs. Moreover, it is the direction, rather than the magnitude, of the perturbation that matters the most. In FGSM the direction is determined by the gradient of the loss function with respect to the model input - x is pushed in the direction of highest loss increase given its true label y .

In the context of entity linking task, we are interested in generating examples adversarial to the learned metric, in other words hard negative and hard positive examples for a given mention m . To this end, we applied the following perturbations to the entity encoder input embeddings $z = \text{input_embed}(e)$:

$$z_{adv}^- = z^- + \epsilon * \text{sign}(\nabla_{z^-} s(m, e^-)) \quad (6)$$

$$z_{adv}^+ = z^+ - \epsilon * \text{sign}(\nabla_{z^+} s(m, e^+)) \quad (7)$$

where m is the anchor mention and z^-, z^+ are the encoder input embeddings of negative and positive entities e^-, e^+ correspondingly.

Given N negative entities for a mention m , the generated adversarial entity embeddings $P_{adv} = \{z_{adv_1}^-, \dots, z_{adv_N}^-, z_{adv}^+\}$ are used as additional training examples, giving rise to an auxiliary loss term that encourages the model to be invariant to local adversarial perturbations. Thus, the final objective we are trying to minimise becomes:

$$L_{Pb}(m, P) + \lambda L_{Pb}(m, P_{adv}) \quad (8)$$

where λ is a hyperparameter controlling the relative contributions of the two losses.

5 Experiments

5.1 Datasets

MedMentions This is a biomedical entity-linking dataset consisting of over 4,000 PubMed abstracts (Mohan and Li, 2019). As recommended by the authors, we use the ST21PV subset, which has around 200,000 mentions in total. A large number of mentions in both the validation and test splits are zero-shot, meaning their ground truth label is not present in the training data. We do not carry out any additional preprocessing on the dataset. Finally,

	MedMentions			Zero-Shot EL		
	Train	Val	Test	Train	Val	Test
Mentions	120K	40K	40K	49K	10K	10K
Entities	19K	8K	8K	333K	90K	70K
% Entities seen	100	57.5	57.5	100	0	0

Table 1: Statistics of datasets used. "% Entities seen" signifies the percentage of ground truth entities seen during training.

for the knowledge base (KB), we follow the framework in Varma et al. (2021) and use the UMLS 2017AA version filtered by the types present in the ST21PV subset. The final KB includes approximately 2.36M entities.

To evaluate our models in the NIL detection setting, we have created a new dataset based on MedMentions. In this dataset, we have assigned mentions corresponding to 11 entity types a NIL label and removed them from the Knowledge Base. Details on the dataset statistics and removed entity types can be found in the Appendix.

Zero-Shot Entity Linking dataset ZESHEL, a general domain dataset was constructed by Logeswaran et al. (2019) from Wikias¹. It consists of 16 independent Wikias. The task is to link mentions in each document to a Wikia-specific entity dictionary with provided entity descriptions. The dataset is zero-shot, meaning there is no overlap in entities between training, validation and test sets.

5.2 Input Representation and Model Architecture

Similarly to Wu et al. (2020); Zhang and Stratos (2021); Varma et al. (2021) our candidate retriever is a bi-encoder consisting of two independent BERT transformers. We use the bi-encoder to encode a textual mention and an entity description independently then obtain a similarity score between them.

Namely, Given a mention and its surrounding context τ_m and an entity τ_e , we obtain dense vector representations $\mathbf{y}_m = \text{red}(T_1(\tau_m))$ and $\mathbf{y}_e = \text{red}(T_2(\tau_e))$, where T_1 and T_2 are the two independent transformers of the bi-encoder and $\text{red}(\cdot)$ is a function that reduces the output of a transformer into a single vector. We use a mean pooling operation for the function $\text{red}(\cdot)$.

As in Wu et al. (2020); Zhang and Stratos (2021); Varma et al. (2021) we use the dot product to score the mention \mathbf{y}_m against an entity vector \mathbf{y}_e when

¹<https://wikia.com>

using the CE loss. For our Proxy-based loss we use cosine similarity.

In this work, we focus on entity linking by efficient candidate retrieval, but we also include the ranker results using the highest scoring candidate entities in the Appendix, where we also include more details on entity, mention and context modelling.

5.3 Training & Evaluation Details

In all our experiments we used the transformer architecture (Vaswani et al., 2017) for the encoders. Namely, we used BERT (Devlin et al., 2019), initialised with appropriate pre-trained weights: SapBERT (Liu et al., 2021) for MedMentions and the uncased BERT-base (Devlin et al., 2019) for ZESHEL. For FGSM regularization, we apply adversarial perturbations to the composite token embeddings (i.e. sum of word, position and segment embeddings) used as input to BERT. We apply our regularization to both Proxy-based and CE. For information on hyperparameter tuning please refer to the Appendix. We tune all of our experiments on the validation set and report results on the test set. Due to hardware limitations, the training was conducted on a single V100 GPU machine with 16 GB of GPU memory. The limited GPU capacity, in particular, memory, posed a challenge by constraining us to using a relatively low number of negatives when training a retriever.

5.3.1 Candidate Retriever

The retriever model is optimised with the Proxy-based loss (3) and benchmark CE loss (1) for fair comparison. We evaluate the retriever on the micro-averaged recall@1 and recall@64 metrics, where in our setup recall@1 is equivalent to accuracy. Here we focus on the recall@1 metric, which is highly relevant for efficient candidate retrieval models that do not necessitate running an expensive cross-encoder for candidate ranking. We use two negative sampling techniques: (1) Random, where the negatives are sampled uniformly at random from all entities in the knowledge base, and (2) Mixed-p: p percent of the negatives are hard, the rest are random. This is motivated by the results shown in Zhang and Stratos (2021). We set the p to 50%.

Hard negative mining Retrieving hard negatives requires running the model in the inference mode over the entire KB. Then, for each mention, the

most similar (i.e. hard) negatives are sampled according to a scoring function. Here, we use FAISS (Johnson et al., 2019) for obtaining hard negatives given a mention and an index of entity embeddings from the KB.

Running a forward pass over the entire KB at regular intervals can be costly and time-consuming as the KB often amounts to millions of entities. Moreover, the computational complexity of retrieving hard negatives may grow exponentially depending on the scoring function. For example, the traditionally used scoring function also leveraged in this work, where the mention and entity are both represented with a single embedding requires $\mathcal{O}(me)$ approximate nearest neighbour searches, where m and e are the number of mentions and entities respectively. However, employing an alternative scoring function such as the *sum-of-max* used in Zhang and Stratos (2021) which requires comparing a set of mention embeddings with a set of entity embeddings results in $\mathcal{O}(mexy)$ where x and y is the number of mention vector and entity vector embeddings. In Zhang and Stratos (2021) x and y are set to 128, the number of maximum tokens in the mention and entity input sequence.

This highlights the computational cost of hard negative mining and underlines the need for both methods which can work effectively with random samples as well as more efficient hard negative mining strategies. In this work we propose a method for the former.

Biomedical Out of Knowledge Base Detection

For the biomedical NIL detection scenario training proceeds exactly as in the in-KB setting. We train models with the Proxy-based loss with different margins, and also a model with the CE loss. In each case, we use a validation set that includes NIL mentions to select an appropriate threshold for the retrieval model. Mentions whose corresponding top-ranked entity does not achieve this score are assigned the NIL label. We choose the threshold that maximises the F1 score for NIL entities in the validation set. We then apply this threshold to detect NIL mentions in the test set.

6 Results

We present the results for candidate retrieval and benchmark our models against suitable methods. We name our method Proxy-based Entity Linking (PEL-Pb). We also report the results of a version of our model which uses the CE (PEL-CE) loss on

	# Neg.	recall@1	recall@64
Angell et al. (2021)	-	50.8	85.3
Agarwal et al. (2021)	-	72.3	95.6
Varma et al. (2021)	100	71.7	-
PEL-CE	32 (mixed)	72.1	95.5
	64 (mixed)	72.1	95.6
	64 (random)	55.7	94.0
PEL-Pb	32 (mixed)	71.6	93.3
	64 (mixed)	72.6	95.0
	64 (random)	63.3	95.9
PEL-CE + FGSM	32 (mixed)	72.3	95.5
PEL-Pb + FGSM	32 (mixed)	72.4	93.7

Table 2: Candidate retrieval results on the MedMentions dataset. CE and Pb refers to cross-entropy and proxy-based losses respectively. All experiments were run with mixed random and hard negatives “(mixed)”, or only “(random)” negatives. The bold figures represent the best score for each recall metric. Note that FGSM PEL variants were only run with 32 negatives due to GPU memory constraints.

	Random		Mixed	
	recall@1	recall@64	recall@1	recall@64
Wu et al. (2019)*	-	81.80	46.5	84.8
Agarwal et al. (2021)	38.6	84.0	50.4	85.1
Ma et al. (2021)	45.4	90.8	-	-
Zhang and Stratos (2021)	-	87.62	-	89.6
PEL-CE	44.1	84.8	52.5	87.2
PEL-Pb	48.9	85.2	53.1	86.0
PEL-CE + FGSM	44.1	85.2	53.2	87.2
PEL-Pb + FGSM	49.7	85.6	54.2	86.6

Table 3: Candidate retrieval results on the ZESHEL dataset. CE and Pb refers to cross-entropy and proxy-based losses respectively. The negative to positive sample ratio for all PEL runs is 32. The bold figures represent the best score for each sampling strategy (random vs. mixed random and hard). The highlighted figure represents the best overall score across strategies. * we use the results reported in Zhang and Stratos (2021) for random negatives and Ma et al. (2021) for mixed negatives.

all experiments for comparison.

6.1 MedMentions

Table 2 shows that all approaches using bi-encoder transformer models strongly outperform the N-Gram TF-IDF proposed in Angell et al. (2021) for recall@1 and also recall@64. We also observe the strong positive effect of including hard negatives during model training. The effect is particularly strong for the CE loss, where recall@1 increases by 17% compared with training on random negatives. We believe that such difference is partly due to the large size of the KB MedMentions KB, amounting to 2.36M entities, which contributes to the impor-

tance of hard negative mining. For the Proxy-based loss, including hard negatives increases recall@1 by 9%, achieving state-of-the-art performance of 72.6%. Adding FGSM regularisation boosted performance, as can be seen from the experiments with 32 negatives (the largest number of negatives we could fit into GPU memory when applying FGSM). However, it did not exceed the performance of the unregularized model with 64 negative samples.

	NIL			All classes incl. NIL	
	auPR	Precision	Recall	Recall@1	Recall@64
Pb (m=0)	83.7	81.2	71.0	72.6	90.4
Pb (m=0.01)	84.4	81.6	71.5	72.5	90.2
Pb (m=0.05)	85.8	83.3	73.5	72.4	89.9
Pb (m=0.1)	87.6	85.2	79.2	69.4	85.7
CE	32.3	31.8	74.0	64.4	76.1

Table 4: NIL detection results on the MedMentions dataset. auPR, precision and recall are reported exclusively for the NIL class, whereas micro-averaged recall@1 and recall@64 are reported for all classes including NIL. Pb: Proxy-based with margin m , CE: Cross-Entropy.




Figure 2: Precision Recall curves for NIL detection on the MedMentions dataset. Pb: Proxy-based with margin m , CE: Cross-Entropy.

Biomedical Out of Knowledge Base detection

We also evaluated our proposed loss function on NIL detection. All models trained with the Proxy-based loss significantly outperform the CE-based model in terms of both precision and recall (Figure 2). The CE loss does not encourage low scores in absolute value for negatives examples, but rather encourages scores that are lower than the scores of positive examples. As we can see from the results, CE training fails to assign low scores to NIL mentions, as these are out-of-distribution negatives and thus have not been compared to positive examples during model training. Our Proxy-based loss does not suffer from this issue, even with a margin of 0.

We believe that this is accomplished by the decoupling of the positive and negative loss terms, such that low *absolute* score values are encouraged for negative examples.

Furthermore, the higher the Proxy-based margin the better the model’s performance with respect to detecting NIL mentions. At the same time, Proxy-based models with lower margins perform better at the overall recall metrics (Figure 2). These metrics are computed with respect to all classes including the NIL class. Given that the performance differences among models with different margins are minimal, a practitioner could choose how to set the margin considering the trade-off between NIL detection and overall model performance. To our knowledge, we are the first to propose a method for NIL detection using the bi-encoder architecture.

6.2 Zero-Shot Entity Linking dataset

Based on the candidate retrieval results in Table 3, we can conclude six key points. (1) Proxy-based models (Pb) outperform their Cross-Entropy (CE) counterparts across all considered settings for recall@1. In particular, our Proxy-based model using hard negatives and FGSM regularization achieves state-of-the-art recall@1 on this dataset. This highlights the gain that we get by breaking the dependency between positive and negative pairs. (2) Including hard negatives always boosts model performance. This is particularly evident on the recall@1 metric. The model trained with CE loss strongly depends on hard negatives, with recall@1 increasing by 8% compared to training with random negatives. For the Proxy-based loss the increase is 4%, as the model already performs competitively when trained with random negatives. This showcases the importance of hard negative sampling for the CE loss. Hard negatives provide the model with much more meaningful feedback and avoid the threat of vanishing gradients (Eq. 2). (3) The difference between Pb and CE models becomes much smaller for recall@64. Trivially, as k increases, recall@ k for all models will converge towards 1. Additionally, as k increases to above the number of hard negatives, the model’s ability to distinguish the hard negatives from the positive will not be seen in the metric. (4) CE models marginally outperform Pb models with hard negatives at recall@64. Hard negatives consistently have a larger impact on CE compared to Pb also at recall@64 (2), while the benefits of Pb have been nullified as discussed in

(3). (5) Alternative methods leveraging the CE loss and different model architectures such as MuVER (Ma et al., 2021) and SOM (Zhang and Stratos, 2021) outperform the bi-encoder based approach at recall@64. However, both MuVER and SOM are more complex models tuned for achieving high recall@64, whereas the main focus of our approach is high recall@1 in the pursuit of avoiding the additional ranking stage. Pb outperforms the only single stage entity linking model Agarwal et al. (2021) across the board. (6) FGSM regularization boosts the results of both Proxy-based and CE models, demonstrating its promise as a general method for regularizing the retrieval model.

7 Discussion & Future Work

We have proposed and evaluated a novel proxy-based loss for biomedical candidate retrieval. Additionally, we have adopted an adversarial regularization technique designed to simulate hard negatives, and shown that both our loss and regularization boost performance on the recall@1 metric. We have also constructed a biomedical dataset for NIL detection and demonstrated that our candidate retrieval model can robustly identify biomedical NIL entities, while maintaining high overall performance. These are important advances towards closing the gap between the two-stage approach that include an expensive cross-encoder and a candidate retriever-only setup.

Notably, our work highlights the importance of hard negative sampling when optimising the candidate generator with the CE loss. Random negative sampling together with CE loss can result in the problem becoming trivial, for example the randomly sampled negative entity having a different type. However, accessing hard negative examples during model training can be challenging, particularly when the knowledge base is large and entity representations are frequently updated.

Considering this, we recommend to employ our Proxy-based loss for the candidate retrieval task in three different scenarios: (1) training with random negatives, (2) optimising for recall@1, (3) detecting NIL entities. Moreover, we also recommend leveraging FGSM regularisation in any setup and both retrieval and ranking tasks.

An interesting approach would be to attempt to approximate hard negatives without frequent updates of the entity representations. This could potentially be done by keeping the entity encoder

frozen, or exploring alternative relatedness measures which does not require frequently running the model over the whole knowledge base. Finally, there is a plethora of work on proxy-based (Movshovitz-Attias et al., 2017; Kim et al., 2020) and pair-based losses (Bromley et al., 1993; Chopra et al., 2005; Schroff et al., 2015; Dong and Shen, 2018), usually discussed in the computer vision and metric learning literature. Improving the candidate retrieval is a crucial step towards high-performing and efficient entity linking systems that can be easily applied in real-world settings.

Limitations

There are several limitations of our work. Firstly, we only demonstrate the advantages of our proposed method when computing hard negatives is computationally expensive, which is the case with large knowledge bases and expensive scoring methods. If computing hard negatives is not a bottleneck, one may use negative sampling with the baseline CE loss. However, biomedical knowledge bases typically contain a huge number of entities. Secondly, in our experiments we were limited to single GPU machines with at most 16GB of GPU memory. This prevented us from including more than 64 negatives samples in the standard setup and 32 negative samples when using FGSM regularization, which could potentially be benefit model performance. Thirdly, we acknowledge that some comparison to related work is missing, in particular, Zhang and Stratos (2021). We were not able to reproduce the results cited in the paper using the publicly available code. Finally, our work is limited to proxy-based metric learning losses. More space could be devoted to the topic of how one could utilise metric learning more broadly for biomedical entity linking. We leave this for future work.

Ethics Statement

The BERT-based models fine-tuned in this work and datasets are publicly available. We will also make our code as well as the biomedical out of knowledge base detection dataset publicly available.

The task of entity linking is often crucial for downstream applications, such as relation extraction, hence potential biases at the entity linking stage can have significant harmful downstream consequences. One source of such biases are the pre-trained language models fine-tuned in this work.

There is a considerable body of work devoted to the topic of biases in language models. One way the entity linking systems can be particularly harmful is when they commit or propagate errors in the language models, knowledge bases, mention detection across certain populations such as races or genders. Because of the high ambiguity across biomedical mentions and entities in the knowledge base, it is important that the users investigate the output prediction of the entity linking system and often take a suggestion, rather than gold standard. Finally, we highlight that linking the entity to its entry in the knowledge base and out of knowledge base detection can be analogous to surveillance and tracking in the computer vision domain, which comes with substantial ethical considerations.

Acknowledgements

We thank Dane Corneil, Georgiana Neculae and Juha Iso-Sipilä for helpful feedbacks and the anonymous reviewers for constructive comments on the manuscript.

References

- Dhruv Agarwal, Rico Angell, Nicholas Monath, and Andrew McCallum. 2021. Entity linking and discovery via arborescence-based supervised clustering. *arXiv preprint arXiv:2109.01242*.
- Rico Angell, Nicholas Monath, Sunil Mohan, Nishant Yadav, and Andrew McCallum. 2021. [Clustering-based inference for biomedical entity linking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2598–2608, Online. Association for Computational Linguistics.
- Rajarshi Bhowmik, Karl Stratos, and Gerard de Melo. 2021. Fast and effective biomedical entity linking using a dual encoder. *ArXiv*, abs/2103.05028.
- Olivier Bodenreider. 2004. [The Unified Medical Language System \(UMLS\): integrating biomedical terminology](#). *Nucleic acids research*, 32(Database issue):D267–D270.
- Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature verification using a "siamese" time delay neural network. In *Proceedings of the 6th International Conference on Neural Information Processing Systems, NIPS'93*, page 737–744, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Razvan Bunescu and Marius Pasca. 2006. Using encyclopedic knowledge for named entity disambiguation.

- Association for Computational Linguistics, European Chapter.
- S. Chopra, R. Hadsell, and Y. LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546 vol. 1.
- 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xingping Dong and Jianbing Shen. 2018. Triplet loss in siamese network for object tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldridge, Eugene Ie, and Diego Garcia-Olano. 2019. [Learning dense representations for entity retrieval](#). *CoNLL 2019 - 23rd Conference on Computational Natural Language Learning, Proceedings of the Conference*, pages 528–537.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*.
- Swapna Gottipati and Jing Jiang. 2011. Linking entities to a knowledge base with query expansion. Association for Computational Linguistics.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and J. Weston. 2019. Poly-encoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring. *arXiv: Computation and Language*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Omar Khattab and Matei Zaharia. 2020. [ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT](#), page 39–48. Association for Computing Machinery, New York, NY, USA.
- Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. 2020. Proxy anchor loss for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-end neural entity linking. *arXiv preprint arXiv:1808.07699*.
- Nevena Lazic, Amarnag Subramanya, Michael Ringgaard, and Fernando Pereira. 2015. Plato: A selective context model for entity resolution. *Transactions of the Association for Computational Linguistics*, 3:503–515.
- 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*, pages 4228–4238, Online. Association for Computational Linguistics.
- Lajanugen Logeswaran, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, Jacob Devlin, and Honglak Lee. 2019. [Zero-shot entity linking by reading entity descriptions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3449–3460, Florence, Italy. Association for Computational Linguistics.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. [Sparse, dense, and attentional representations for text retrieval](#). *Transactions of the Association for Computational Linguistics*, 9:329–345.
- Xinyin Ma, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Weiming Lu. 2021. [MuVER: Improving first-stage entity retrieval with multi-view entity representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2617–2624, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Takeru Miyato, Andrew M Dai, and Ian Goodfellow. 2016. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*.
- Sunil Mohan and Donghui Li. 2019. [Medmentions: A large biomedical corpus annotated with UMLS concepts](#). *CoRR*, abs/1902.09476.
- Jose G Moreno, Romaric Besançon, Romain Beaumont, Eva D’hondt, Anne-Laure Ligozat, Sophie Rosset, Xavier Tannier, and Brigitte Grau. 2017. Combining word and entity embeddings for entity linking. In *European Semantic Web Conference*, pages 337–352. Springer.
- Yair Movshovitz-Attias, Alexander Toshev, Thomas K. Leung, Sergey Ioffe, and Saurabh Singh. 2017. No fuss distance metric learning using proxies. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Lin Pan, Chung-Wei Hang, Avirup Sil, Saloni Potdar, and Mo Yu. 2021. Improved text classification via contrastive adversarial training. *arXiv preprint arXiv:2107.10137*.

Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. **Facenet: A unified embedding for face recognition and clustering**. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823.

Ozge Sevgili, Artem Shelmanov, Mikhail Arkhipov, Alexander Panchenko, and Chris Biemann. 2020. Neural entity linking: A survey of models based on deep learning. *arXiv preprint arXiv:2006.00575*.

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):443–460.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.

Chen-Tse Tsai and Dan Roth. 2016. Cross-lingual wiki-fication using multilingual embeddings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 589–598.

Maya Varma, Laurel Orr, Sen Wu, Megan Leszczynski, Xiao Ling, and Christopher Ré. 2021. **Cross-domain data integration for named entity disambiguation in biomedical text**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4566–4575, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefinukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.

Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2019. **Zero-shot Entity Linking with Dense Entity Retrieval**.

Ledell Wu, Fabio Petroni, Martin Josifoski, 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)*, pages 6397–6407, Online. Association for Computational Linguistics.

Wenzheng Zhang and Karl Stratos. 2021. **Understanding hard negatives in noise contrastive estimation**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*,

pages 1090–1101, Online. Association for Computational Linguistics.

Appendices

A Context and Mention Modelling

We represent a mention and its surrounding context, τ_m , as a sequence of word piece tokens

[CLS] ctxt_l [M_s] mention [M_e] ctxt_r [SEP]

where mention, ctxt_l and ctxt_r are the word-piece tokens of the mention, left and right context, and [M_s] and [M_e] are special tokens marking the start and end of a mention respectively.

Due to the differences in available data, we represent entities differently for ZESHEL and MedMentions. On ZESHEL, we represent entities with a sequence of word piece tokens

[CLS] title [ENT] description [SEP]

where [ENT] is a special separator token. In contrast, when training on the MedMentions dataset we represent an entity by the sequence

[CLS] title [SEP] types [SEP] description [SEP]

Descriptions of entities were sourced from UMLS.

B Candidate ranker setup and results

To evaluate the impact of our candidate retriever model on the downstream task of candidate ranking, we also conducted ranking experiments on both datasets.

		# Candidates	Ranker	Accuracy
ZESHEL	Wu et al. (2020)	64	Base	61.3
	Wu et al. (2020)	64	Large	63.0
	Zhang and Stratos (2021)	64	Base	66.7
	Zhang and Stratos (2021)	64	Large	67.1
	PEL-Pb	16	Base	62.8
	PEL-Pb + FGSM	16	Base	64.6
MedMentions	Bhowmik et al. (2021)*	-	-	68.4
	Angell et al. (2021)	-	-	72.8
	Varma et al. (2021)	10	Base	74.6
	PEL-Pb	16	Base	74.0
	PEL-Pb + FGSM	16	Base	74.6
	Angell et al. (2021)	-	-	74.1
	+ post-processing	-	-	74.1
	Varma et al. (2021)	10	Base	74.8
+ post-processing	-	-	74.8	

Table 5: Ranker results on the ZESHEL and MedMentions datasets. * uses the full MedMentions dataset, rather than the ST21PV subset used by other models reported in the table and recommended by MedMentions authors’.




Figure 3: Comparison of smoothed gradient norms over training steps using two losses, CE and Proxy-based. The left plot visualizes the smoothed gradient norm when using random, and the right one leveraging mixed-50% negatives. All the experiments were conducted on ZESHEL using 32 negatives.

Training & Evaluation setup Similarly as in related work (Logeswaran et al., 2019; Wu et al., 2020; Zhang and Stratos, 2021), the highest scoring candidate entities from the candidate retriever are passed to a ranker, which is a cross-encoder consisting of one BERT transformer. The cross-encoder Logeswaran et al. (2019) is used to select the best entity out of the candidate pool. It takes as input $\tau_{m,e}$, which is the concatenation of mention/context and entity representations τ_m and τ_e . We then obtain a dense vector representation for a mention-entity pair $\mathbf{y}_{m,e} = T_{\text{cross}}(\tau_{m,e})$, where $T_{\text{cross}}(\tau_{m,e})$ is the BERT transformer of the cross-encoder and $\text{red}(\cdot)$ is a mean pooling function that takes the mean over input tokens embeddings. Entity candidates are scored by applying a linear layer $s_{\text{cross}}(m, e) = \mathbf{y}_{m,e} \mathbf{W}$.

We pick the best performing retrieval model on recall@16 and use it to retrieve top 16 candidate entities for each mention. As the number of candidate entities is relatively low, we do not perform negative sampling and optimise the cross-encoder with the CE loss (Eq. 1). We report the micro-averaged unnormalized accuracy on the MedMentions dataset and macro-averaged unnormalized accuracy on the ZESHEL dataset in line with the prior work (Zhang and Stratos, 2021; Wu et al., 2020). The results are shown in the Table 5.

Results In Table 5 we can observe the downstream effect of having a candidate generator model with high recall@1 performance. On ZESHEL, We can see that a cross-encoder trained with the top 16 candidates from our best performing candidate generator achieved higher accuracy than Wu et al. (2020) who used the top 64 candidates. Moreover, similarly as with the candidate retrieval,

FGSM boosts performance. For completeness, we have also included the state-of-the-art results from Zhang and Stratos (2021) who used 64 candidates and a larger BERT model in the cross-encoder. In our experiments we were limited to a single GPU with 16 GB memory which restricted us to a low number of maximum candidates, namely 16. We strongly believe that including more candidates than 16 would boost the performance of our method.

On MedMentions a cross-encoder trained with the top 16 candidates from our best performing candidate generator model achieved a competitive accuracy of 74%. The accuracy further increased to 74.6% when adding FGSM regularisation, coming close to the state-of-the-art performance of Varma et al. (2021), which includes additional post-processing.

C Training details

The hyperparameters used for conducting the experiments are visible in Table 6. We use a single NVIDIA V100 GPU with 16 GB of GPU memory for all model trainings.

D Biomedical Out of Knowledge Base dataset details

We constructed the OKB dataset by replacing the label of a set of mentions from the MedMentions corpus (Mohan and Li, 2019) with the NIL class. Namely we pick the mentions belonging to 11 types: *Mental Process, Health Care Related Organization, Element Ion or Isotope, Medical Device, Health Care Activity, Diagnostic Procedure, Professional or Occupational Group, Mental Process, Laboratory Procedure, Regulation or Law,*

Param	Bi-encoder	Cross-Encoder
Input sequence length	128	256
learning rate	1e-5	2e-5
warmup proportion	0.25	0.2
eps	1e-6	1e-6
gradient clipping value	1.0	1.0
effective batch size	32	4
epochs	7	5
learning rate scheduler	linear	linear
optimiser	AdamW	AdamW
α	32	-
δ	0.0	-
FGSM λ	1	1
FGSM ϵ	0.01	0.01

Table 6: Learning parameters for the bi-encoder and cross-encoder.

Organization, Professional Society. The final OKB subset includes approximately 24K mentions and 3K unique entities.

To ensure that the OKB dataset does not suffer from easy inferences and allows us to evaluate model performance. We ensured that the zero-shot distribution of the OKB mentions and types across the train/validation/test split was in line with the zero-shot distribution of mentions and types in the whole dataset. Additionally, we verified that there is no significant overlap between mention surface forms across the splits. Moreover, we looked at the length of entity descriptions which are used to create entity representations checking that the OKB mentions entity representations statistics are similar to the statistics computed using the whole dataset.

E Gradient norm analysis

	Train	Dev	Test
Mentions	14K	4.8K	4.7K
Entities	2.2K	1.1K	1.1K
% Entities seen	100	57.7	57.5

Table 7: Statistics of the OKB MedMentions subset.

Figure 3 shows the behaviour of the gradient l_2 norm for both losses. We can see that for both random and mixed negatives, the norm of the Proxy-based loss has considerably lower variance. This is visible particularly when using the mixed negatives.