

# Find the Funding: Entity Linking with Incomplete Funding Knowledge Bases

**Gizem Aydin**

Radboud University

gizemaydin96@gmail.com

**Seyed Amin Tabatabaei**

Elsevier

s.tabatabaei@elsevier.com

**Giorgios Tsatsaronis**

Elsevier

g.tsatsaronis@elsevier.com

**Faegheh Hasibi**

Radboud University

f.hasibi@cs.ru.nl

## Abstract

Automatic extraction of funding information from academic articles adds significant value to industry and research communities, including tracking research outcomes by funding organizations, profiling researchers and universities based on the received funding, and supporting open access policies. Two major challenges of identifying and linking funding entities are: (i) sparse graph structure of the Knowledge Base (KB), which makes the commonly used graph-based entity linking approaches suboptimal for the funding domain, (ii) missing entities in KB, which (unlike recent zero-shot approaches) requires marking entity mentions without KB entries as NIL. We propose an entity linking model that can perform NIL prediction and overcome data scarcity issues in a time and data-efficient manner. Our model builds on a transformer-based mention detection and a bi-encoder model to perform entity linking. We show that our model outperforms strong existing baselines.

## 1 Introduction

Entity Linking (EL) aims to annotate text with corresponding entity identifiers from a Knowledge Base (KB) and is a building block for different tasks, such as document ranking (Xiong et al., 2017), entity retrieval (Hasibi et al., 2016), and question understanding in conversations (Shang et al., 2021). Recent years have witnessed the flourishing of entity linking approaches for zero-shot (Wu et al., 2020; Li et al., 2020) and open-domain setups (van Hulst et al., 2020; Cao et al., 2021). While zero-shot entity linking can generalize to new specialized domains and entity dictionaries, existing approaches cannot perform NIL prediction; i.e., identifying entity mentions without a target entity in a knowledge base and assigning them to NIL. Open-domain entity linkers, on the other hand, build on the availability of rich entity relations and descriptions in KBs. This makes ex-

isting EL approaches suboptimal for real-world applications of entity linking in domains with incomplete knowledge bases, where both in-KB and out-of-KB entities should be identified.

In this paper, we aim to address entity linking in the funding domain (Dai et al., 2021; Alexander and de Vries, 2021), which is essential for funding organizations to track the outcome of the research they funded (Kayal et al., 2019) and also helps to comply with open access rules (Dai et al., 2021). Knowledge bases of funding organizations, either proprietary or open access (e.g., the funding KB Crossref<sup>1</sup>), contain brief information about entities (e.g., official name and acronym). They also have extremely sparse graph structure with large amount of missing entities that need to be found from research articles. This implies that EL in the funding domain requires detecting mentions with out-of-KB entities while handling sparse entity relations and descriptions. The approach, should also be able to operate with limited training data, as large public datasets are rarely available for domain-specific applications.

We propose a two-step EL approach, where we first identify entity mentions using task adaptive pre-training (Gururangan et al., 2020) of BERT (Devlin et al., 2019) and then perform Entity Disambiguation (ED) by utilizing a bi-encoder model to learn dense entity and mention representations. Our bi-encoder model and training approach using negative sampling are specifically designed to operate with out-of-KB entities. The bi-encoder is followed by a modest feature-based model to map the entities to an entity in KB or NIL. We create two new datasets for EL in the funding domain and compare our mention detection and ED approaches with strong neural and feature-based models. We show that our model improves over existing baselines for both entity disambiguation and end-to-end entity linking.

<sup>1</sup><https://www.crossref.org>

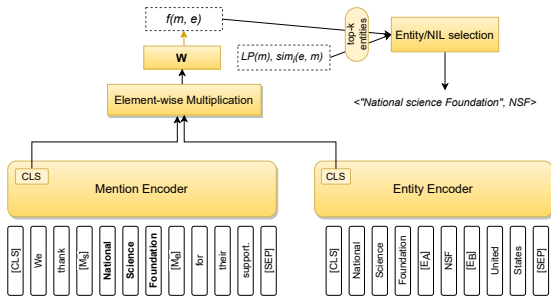


Figure 1: Overview of our entity disambiguation approach for incomplete proprietary knowledge bases.

In summary, our contributions include: (i) proposing a data-efficient model for entity linking (with NIL prediction) in funding domain that is efficient and can be used with modest computational power, (ii) improving upon existing EL approaches for funding organization, and (iii) releasing new training and evaluation datasets for entity linking in funding domain. To our knowledge, this is the first and largest publicly available dataset for entity linking in funding domain. The code and datasets created in this paper are made publicly available.<sup>2</sup>

## 2 Method

In this section, we provide a formal definition of the task, followed by the description of our Mention Detection (MD) approach and Funding entity Disambiguation model, referred to as FunD.

### 2.1 Task Definition

We denote  $\mathcal{E}$  as the set of entities in a knowledge base, where each entity  $e \in \mathcal{E}$  is accompanied by a textual description. Let  $m = (s, t) \in M$  denote an entity mention with start and end positions  $s$  and  $t$ . Given a document  $d = \{w_1, w_2, \dots, w_n\}$ , our aim is to generate the list:

$$\mathcal{L} = \{(\langle s, t \rangle, a) \mid 1 \leq s \leq t \leq n, a \in \mathcal{E} \cup \text{NIL}\},$$

which represents all possible mentions linked to an entity in the KB (*in-KB setup*) or NIL (*out-of-KB setup*). This task is similar to zero-shot (Wu et al., 2020) and open domain entity linking (van Hulst et al., 2020), but different from them, entities do not need to have an entry in the KB.

### 2.2 Mention Detection

For mention detection, we adapt BERT (Devlin et al., 2019) to the funding domain. Domain Adaptive Pre-Training (DAPT), while being effective,

requires a large amount of domain-specific text, requires large amount of training data which is not feasible for the funding domain (Gerritse et al., 2022; Nogueira et al., 2019). We therefore utilize the Task-Adaptive Pre-Training (TAPT) (Gururangan et al., 2020), which requires a far smaller but more task-relevant training corpus and is proven to be more effective than DAPT. We train BERT with the Masked Language Model objective on acknowledgments of research papers. We refer to this model as BERT<sub>TAPT</sub>. We then fine-tune BERT<sub>TAPT</sub> for the mention detection task using IOB tags.

### 2.3 Entity Disambiguation

**Candidate Entity Selection** To obtain the likelihood of an entity being a target link of a mention, we employ a bi-encoder model (Wu et al., 2020) for encoding a mention (with its context) and an entity. Our bi-encoder model utilizes two BERT encoders for generating entity and mention representations. The entity encoder takes the structured entity description as the input:

$$x_e = \text{BERT}_{[\text{CLS}]}([\text{CLS}] \text{ val}_{A_1} [\text{E}_A] \dots [\text{E}_A] \text{ val}_{A_n} [\text{E}_B] \text{ val}_B [\text{SEP}]),$$

where  $[\text{E}_A]$  and  $[\text{E}_B]$  are two word-piece tokens, selected among the unused tokens of BERT, and  $\text{val}_{A_i}$  and  $\text{val}_B$  denote values for entity attributes  $A$  and  $B$ . Here  $A$  corresponds to names of entities, which is a multi-valued attribute, and  $B$  is the country of the funding organization. For the mention encoder, we follow Wu et al. (2020) and obtain mention representations by:

$$x_m = \text{BERT}_{[\text{CLS}]}([\text{CLS}] \text{ ctxt}_{\text{left}} [\text{M}_s] \text{ mention} [\text{M}_e] \text{ ctxt}_{\text{right}} [\text{SEP}]),$$

where  $\text{ctxt}_{\text{left}}$  and  $\text{ctxt}_{\text{right}}$  represent context words before and after the mention. In this work, both BERT models are initialized with BERT<sub>TAPT</sub>. The mention-entity score is then obtained by:

$$f(m, e) = \mathbf{W}^T(x_m \odot x_e), \quad (1)$$

where  $\odot$  refers to element-wise multiplication of mention representation  $x_m$  and entity representation  $x_e$ , and  $\mathbf{W} \in \mathcal{R}^{\text{BERT} \times 2}$  represents learnable weights. The binary cross-entropy loss  $L$  is used to train the model:

$$L = -\frac{1}{N} \sum_{i=1}^N l_i \log(f(m_i, e_i)) - (1 - l_i) \log(1 - f(m_i, e_i)),$$

<sup>2</sup><https://github.com/informagi/Fund-EL>

where  $N$  is the number of training examples and  $l_i$  is a binary value that is set to 1 if  $e_i$  is the correct entity for mention  $m_i$ .

**Negative Sampling** Following Gillick et al. (2019), we perform training in rounds, where the model obtained in each round is used to produce hard negatives for the next round. Contrary to (Gillick et al., 2019), we do not use in-batch random negative sampling, as it provides less diverse random negatives for sparse domain-specific applications compared to open-domain EL.

The following strategy is employed for random and hard negative sampling. In the first round, negative entities of each mention are sampled randomly from the entire KB. For the next rounds, both random and hard negatives are used. Hard negatives are entities ranked above the correct entity by the model learned in the previous step. For mentions with out-of-KB entities, the top-K entities are selected as hard negatives (K is set to 10 following (Gillick et al., 2019)). The number of random negatives for each mention is computed based on the number of hard negatives:

$$Neg_r(m) = \lfloor \frac{\sum_{i=1}^{|M|} Neg_h(m)}{|M|} \rfloor, \quad (2)$$

where  $Neg_r()$  and  $Neg_h()$  give number of random and hard negatives, respectively. Using this strategy, we strive a balance between random and hard negatives, while giving hard mentions (i.e., mentions that their correct entities are in low ranks) a larger number of hard negatives.

**Entity or NIL Selection** Once we have obtained candidate entities from our bi-encoder model, we turn to mapping each mention to an entity in the knowledge base or NIL. We employ a feature-based model using Gradient Boosting Machine (GBM) (Friedman, 2001). Our model utilizes five light-weight features: (i) score obtained by the bi-encoder model, (ii) maximum Levenshtein similarity between the mention and the labels of the candidate entity, (iii) link probability of the mention (Balog, 2018) obtained by dividing numbers of times a mention appears as a link by total number of occurrences of a term:  $P(link|m) = n_{link}(m)/n(m)$ , and (iv) commonness score (Balog, 2018) obtained by  $P(e|m) = n(m, e) / \sum_{e' \in \mathcal{E}} n(m, e')$ , with  $n(m, e)$  denoting number of times that mention  $m$  is linked to entity  $e$ . A mention is linked to the entity with the highest GBM score if higher than threshold  $\tau$ .

System	Set	P	R	F1
Stanford NER	Test	73.70	75.10	74.39
Flair <sup>NER</sup>	Test	<b>85.83</b>	78.02	81.74
BERT <sup>NER</sup>	Test	79.18	86.03	82.46
BERT <sup>MD</sup> <sub>TAPT</sub>	Test	80.28	<b>86.54</b>	<b>83.29</b>
Stanford NER	Eval	76.17	72.87	74.48
BERT <sup>MD</sup> <sub>TAPT</sub>	Eval	<b>79.08</b>	<b>85.31</b>	<b>82.08</b>

Table 1: Mention detection results on ELFund dataset.

### 3 Experiments

**Data** We use the Crossref funding registry as our KB, containing information about 25,859 funding organizations. We create two new datasets for the funding domain, ELFund and EDFund (Afzal et al., 2022), which are used for experiments. The datasets are split into training, validation, test, and eval sets, with no overlaps in the training, test, and eval across the two datasets. The validation set is used for searching hyper-parameters. The eval set contains completely unseen production data, used for the final evaluation of the models; see Appendix A and B for more details about the datasets and experimental setup.

**Evaluation Metrics** To evaluate the mention detection step, we use strong matching precision, recall, and F1 score (Tjong Kim Sang and De Meulder, 2003; Usbeck et al., 2019). The ED and EL tasks are evaluated in three settings: (i) *In-KB*, for mentions linked to entities in KB, (ii) *Emerging Entities (EE)* for Out-of-KB entities; i.e., mentions linked to no entities, and (iii) *All*, for in-KB and emerging entities. We evaluate the ED task using micro and macro averaged accuracy for the All setting (Hoffart et al., 2014). For In-KB and EE settings, we report on micro- and macro- averaged precision, recall, and F1 (Usbeck et al., 2019).

### 4 Results

**Mention Detection** Table 1 shows the results for the mention detection step. We compare BERT<sup>MD</sup><sub>TAPT</sub> with Stanford NER (Finkel et al., 2005), Flair<sup>NER</sup> (Akbik et al., 2018), and BERT<sup>NER</sup> (Devlin et al., 2019). The results show that Flair<sup>NER</sup> achieves the highest precision, but the lowest recall compared to the BERT-based models. We also observe that BERT<sup>MD</sup><sub>TAPT</sub> outperforms all baselines with respect to the F1 score, showing the importance of task adaptive pre-training when limited data is available.

Method	Set	All		EE			In-KB		
		Acc <sub>mic</sub>	Acc <sub>mac</sub>	P <sub>mic</sub>	R <sub>mic</sub>	F1 <sub>mic</sub>	P <sub>mic</sub>	R <sub>mic</sub>	F1 <sub>mic</sub>
Commonness	Test	83.8	85.81	53.55	<b>88.2</b>	66.64	<b>94.22</b>	82.99	88.25
GBM <sub>F26</sub>	Test	91.02	<b>92.84</b>	<b>79.11</b>	78.67	78.89	93.2	<b>93.29</b>	93.25
FunD	Test	<b>91.15</b>	92.76	77.44	81.14	<b>79.25</b>	93.82	92.99	<b>93.40</b>
GBM <sub>F26</sub>	Eval	90.26	90.84	<b>80.03</b>	81.49	80.75	92.69	<b>92.3</b>	92.5
FunD	Eval	<b>90.66</b>	<b>91.11</b>	79.26	<b>85.45</b>	<b>82.24</b>	<b>93.56</b>	91.86	<b>92.7</b>

Table 2: Entity disambiguation results on the EDFund dataset. Best results for each set are marked in bold face.

MD	ED	Setting	F1 <sub>mic</sub>	F1 <sub>mac</sub>
Stanford NER	GBM <sub>F26</sub>	All	68.43	69.34
BERT <sub>TAPT</sub> <sup>MD</sup>	FunD	All	<b>75.81</b>	<b>76.59</b>
Stanford NER	GBM <sub>F26</sub>	EE	43.34	71.01
BERT <sub>TAPT</sub> <sup>MD</sup>	FunD	EE	<b>52.82</b>	<b>73.68</b>
Stanford NER	GBM <sub>F26</sub>	In-KB	77.33	74.73
BERT <sub>TAPT</sub> <sup>MD</sup>	FunD	In-KB	<b>85.14</b>	<b>81.86</b>

Table 3: Entity linking results on the ELFund dataset.

**Entity Disambiguation** Table 2 presents entity disambiguation results. We compare our ED method, FunD with two baselines: (i) Commonness (Hasibi et al., 2015), where each mention is linked to the entity with the highest commonness score if it is greater than zero, (ii) GBM<sub>F26</sub>, which is a strong feature-based GBM model with 26 features, ranging from string similarities (e.g. BM25) to statistical features (e.g., commonness and link probability). We note that state-of-the-art EL methods, such as REL (van Hulst et al., 2020), GENRE (Cao et al., 2021), and BLINK (Wu et al., 2020) cannot be used as baselines, as they rely on data resources that are not available in our KB and also do not address NIL prediction. We, however, implemented BLINK’s bi-encoder with a score threshold, and obtained micro average accuracy of 60.16, which is a far worse performance compared to other baselines. Table 2 results show that FunD strives a balance between precision and recall and can achieve the best results with respect to F1 in both In-KB and EE setups. This observation is also mirrored with respect to accuracy on the Eval set.

**Entity Linking** Putting the pieces together, we show the results of end-to-end entity linking on test set of ELFund dataset in Table 3. We compare our model with the best ED baseline (GBM<sub>F26</sub>) combined with Stanford NER (a fast and strong existing MD model). The results indicate that our MD and ED models improve the existing feature based model by a large margin, reinforcing our pre-

System	With GPU	Without GPU
FunD	9.26 ± 0.47	23.07 ± 0.78
GBM <sub>F26</sub>	99.2 ± 0.45	99.2 ± 0.45

Table 4: Efficiency of ED models (in seconds).

vious finding that our models can be successfully applied to the funding domain.

**Efficiency** Finally, we measure the run time of ED model by running it on a random sample of 100 sentences with 306 mentions. We pass 12 candidate entities to both GBM<sub>F26</sub> and FunD and measure the run time in seconds. The experiment is repeated 10 times on a machine with an Intel Xeon E-2276M (2.80GHz, 32GB RAM) CPU and an NVIDIA Quadro T1000 GPU with 4GB memory. Table 4 shows that FunD is four times faster than the feature-based GBM<sub>F26</sub> model without GPU. The difference is even larger using GPU, as FunD’s efficiency is increased, while GBM<sub>F26</sub> performance does not change with GPU. The inefficiency of the GBM<sub>F26</sub> model is mostly attributed to the calculation of the hand-crafted features.

## 5 Conclusions

In this paper, we have introduced an entity linking method for funding domain, where the knowledge base has sparse graph structure and limited information is available about entities. The model builds on BERT to perform mention detection, and a bi-encoder model to conduct the entity disambiguation. We compared our method to strong feature-based and zero-shot models and showed that our model can perform NIL assignments and overcome data scarcity issues more efficient and effective than comparable baselines. As future work, we would like to explore the benefit of employing contrastive learning for the highly ambiguous entity mentions, which could provide further robustness to extracting and linking such entities in scientific texts that span across all sciences.

## Acknowledgements

We thank Ramadurai Petchiappan, Georgios Cheirmpas, and Efthymios Tsakonas, for their assistance in preparing the final released datasets.

## References

- Zubair Afzal, Johan Boots, Heber Mc Mahon, Nishant Mintri, Seyedamin Tabatabaei, and George Tsatsaronis. 2022. [Edfund- Elsevier funding entity linking dataset](#).
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *Proceedings of 27th International Conference on Computational Linguistics (COLING)*, pages 1638–1649.
- Daria Alexander and Arjen P. de Vries. 2021. “this research is funded by.”: Named entity recognition of financial information in research papers. In *Proceedings of the 11th International Workshop on Bibliometric-enhanced Information Retrieval co-located with 43rd European Conference on Information Retrieval (ECIR)*, pages 102 - 110.
- Krisztian Balog. 2018. *Entity-oriented search*. Springer Nature.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. Autoregressive entity retrieval. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, pages 3–7.
- Suyang Dai, Yuxia Ding, Zihan Zhang, Wenxuan Zuo, Xiaodi Huang, and Shanfeng Zhu. 2021. Grantextractor: Accurate grant support information extraction from biomedical fulltext based on Bi-LSTM-CRF. *IEEE/Association for Computing Machinery Transactions on Computational Biology and Bioinformatics*, 18(1):205–215.
- 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 (NAACL)*, pages 4171–4186.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by Gibbs sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 363–370.
- Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Emma J. Gerritse, Faegheh Hasibi, and Arjen P. de Vries. 2022. Entity-aware transformers for entity search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR, SIGIR ’22)*, page 1455–1465.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. 2019. Learning dense representations for entity retrieval. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 528–537.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 8342–8360.
- Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2015. Entity linking in queries: Tasks and evaluation. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval (ICTIR)*, pages 171–180.
- Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2016. Exploiting entity linking in queries for entity retrieval. In *Proceedings of the 2016 ACM on International Conference on the Theory of Information Retrieval, ICTIR ’16*, pages 209–218.
- 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 (WWW)*, pages 385–396.
- Subhradeep Kayal, Zubair Afzal, George Tsatsaronis, Marius Doornenbal, Sophia Katrenko, and Michelle Gregory. 2019. A framework to automatically extract funding information from text. In *Machine Learning, Optimization, and Data Science*, pages 317–328.
- Belinda Z. Li, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. 2020. Efficient one-pass end-to-end entity linking for questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6433–6441.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*.
- Mingyue Shang, Tong Wang, Mihail Eric, Jiangning Chen, Jiyang Wang, Matthew Welch, Tiantong Deng, Akshay Grewal, Han Wang, Yue Liu, Yang Liu, and Dilek Hakkani-Tur. 2021. Entity resolution in open-domain conversations. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 26–33.

Set	#Articles <sub>EL</sub>	#Links <sub>EL</sub>	#Links <sub>ED</sub>	NIL <sub>ED</sub> %
Train	22,720	67,671	95,761	18.58%
Val	991	4,333	5,618	15.47%
Test	3,943	16,355	19,765	15.56%
Eval	13,851	37,495	52,378	18.83%

Table 5: Overview of ELFund and EDFund datasets. For each split, number of articles, number of mentions used in training the NER models, number of mention-entity pairs used in training the ED models, and the percentage of NIL mentions among those pairs are shown.

Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142-147.

Ricardo Usbeck, Michael Röder, Michael Hoffmann, Felix Conrads, Jonathan Huthmann, Axel-Cyrille Ngonga Ngomo, Christian Demmler, and Christina Unger. 2019. Benchmarking question answering systems. *Semantic Web*, 10(2):293-304.

Johannes M. van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P. de Vries. 2020. REL: an entity linker standing on the shoulders of giants. In *Proceedings of the 43rd International Association for Computing Machinery SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 2197-2200.

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.

Chenyang Xiong, Jamie Callan, and Tie-Yan Liu. 2017. Word-Entity Duet Representations for Document Ranking. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 763-772.

## A Dataset Statistics

The ELFund and EDFund datasets are created based on scientific articles published before 2017. Expert annotators were asked to identify sentences that contain funding organizations of the research (e.g., X was funded by source Y) and link the organizations to entities in the Crossref KB. ELFund, further contains sentences that could be also automatically identified by a classifier. Both datasets were annotated by two experts to find the boundary of mentions and their corresponding entities if available. Disagreements were resolved by a third annotator, and mentions with out-of-KB entities are annotated with NIL. We note that ELFund is not a subset of EDFund.

## B Training Configuration

We train the case-preserving version of BERT<sub>BASE</sub> with 2M sentences containing funding information to obtain the BERT<sub>TAPT</sub> model. Unless indicated otherwise, the hyper-parameters recommended by Gururangan et al. (2020) are used for training. The training is done on an NVIDIA Tesla K80 GPU with 12 GB of memory with a batch size of 2048 through gradient accumulation and for one epoch (1000 steps). We further fine-tune BERT<sub>TAPT</sub> for the mention detection task on the ELFund dataset. The fine-tuning process is done for 3 epochs with batch size of 8. We refer to this model as BERT<sub>TAPT</sub><sup>MD</sup>. For disambiguation, the bi-encoder model is trained on the EDFund dataset with a learning rate of  $2 \times 10^{-5}$  and batch size of 16. The training is performed in 4 rounds, each round consisting of 2 epochs. In the first round, 3 random negatives are used for each mention. The score threshold  $\tau$  is set to 0.042 using grid search. Following (Wu et al., 2020), the mention and entity representations are limited to 64 and 256 tokens, respectively.