

# Entity Linking in the Job Market Domain

Anonymous ACL submission

## Abstract

In Natural Language Processing, entity linking (EL) has centered around Wikipedia, but yet remains underexplored for the job market domain. Disambiguating skill mentions can help us get insight into the current labor market demands. In this work, we are the first to explore EL in this domain, specifically targeting the linkage of occupational skills to the ESCO taxonomy (le Vrang et al., 2014). Previous efforts linked coarse-grained (full) sentences to a corresponding ESCO skill. In this work, we link more fine-grained span-level mentions of skills. We tune two high-performing neural EL models, a bi-encoder (Wu et al., 2020) and an autoregressive model (Cao et al., 2021), on a synthetically generated mention–skill pair dataset and evaluate them on a human-annotated skill-linking benchmark. Our findings reveal that both models are capable of linking implicit mentions of skills to their correct taxonomy counterparts. Empirically, BLINK outperforms GENRE in strict evaluation, but GENRE performs better in loose evaluation (accuracy@k).<sup>1</sup>

## 1 Introduction

Labor market dynamics, influenced by technological changes, migration, and digitization, have led to the availability of job descriptions (JD) on platforms to attract qualified candidates (Brynjolfsson and McAfee, 2011, 2014; Balog et al., 2012). It is important to extract and link skills to a disambiguated surface form in JDs, allowing us to quantify the current labor market dynamics and determine the demands and needs. We attempt to tackle the problem of *entity linking* (EL) in the job market domain, specifically the linking of fine-grained span-level skill mentions to a specific surface form.

Generally, EL is the task of linking mentions of entities in unstructured text documents to their respective surface-form entities in a knowledge base

<sup>1</sup>The source code can be found at [https://anonymous.4open.science/r/el\\_esco-E629](https://anonymous.4open.science/r/el_esco-E629)

(KB), most commonly Wikipedia (He et al., 2013). Recent models address this problem by producing entity representations from a (sub)set of KB information, e.g., entity descriptions (Logeswaran et al., 2019; Wu et al., 2020), fine-grained entity types (Raiman and Raiman, 2018; Onoe and Durrett, 2020; Ayoola et al., 2022), or generation of the input text autoregressively (Cao et al., 2021, 2022).

For skill linking specifically, we use the European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al., 2014) taxonomy due to its comprehensiveness. Previous work classified spans to its taxonomy code via multi-class classification (Zhang et al., 2022b) without surrounding context and neither the full breadth of ESCO. Gnehm et al. (2022) approaches it as a sequence labeling task, but only uses more coarse-grained ESCO concepts, and not the full taxonomy. Last, others attempt to match the full sentence to their respective taxonomy title (Decorte et al., 2022, 2023; Clavié and Soulié, 2023).

The latter comes with a limitation: The taxonomy title does not indicate which subspan in the sentence it points to, without an exact match. We define this as an *implicit* skill, where mentions (spans) in the sentence do not have an exact string match with a skill in the ESCO taxonomy. The differences can range from single tokens to entire phrases. For example, we can link “being able to work together” to “plan teamwork”.<sup>2</sup> If we know the exact span, this implicit skill can be added to the taxonomy as an alternative choice for the surface skill. As a result, this gives us a more nuanced view of the labor market skill demands. Therefore, we attempt to train models to the linking of both implicit and explicit skill mentions.

**Contributions.** Our findings can be summarized as follows: ① We pose the task of skill linking as an entity linking problem, showing promising

<sup>2</sup><https://t.ly/3VUJG>

	Instances	Unique Titles	UNK
Train	123,619	12,984	14,641
Dev.	480	149	233
Test	1,824	455	813

Table 1: **Data Statistics.** Data distribution of train, dev, and test splits. UNK indicates skills mentions that are not linked to a corresponding taxonomy title.

079 results of successful linking with two entity linking  
080 systems. ② We present a qualitative analysis showing  
081 that the model successfully links implicit skills  
082 to their respective skill surface form in ESCO.

## 083 2 Methodology

084 **Definition.** In EL, we process the input document  
085  $\mathcal{D} = \{w_1, \dots, w_r\}$ , a collection of entity men-  
086 tions denoted as  $\mathcal{MD} = \{m_1, \dots, m_n\}$ , and a KB,  
087 ESCO in our case:  $\mathcal{E} = \{e_1, \dots, e_{13890}, \text{UNK}\}$ .  
088 The objective of an EL model is to generate a list  
089 of mention-entity pairs  $\{(m_i, e_i)\}_{i=1}^n$ , where each  
090 entity  $e$  corresponds to an entry in a KB. We assume  
091 that both the titles and descriptions of the entities  
092 are available, which is a common scenario in EL  
093 research (Ganea and Hofmann, 2017; Logeswaran  
094 et al., 2019; Wu et al., 2020). We also assume that  
095 each mention in the document has a corresponding  
096 valid gold entity present in the knowledge base,  
097 including UNK. This scenario is typically referred  
098 to as “in-KB evaluation”. Similar to prior research  
099 efforts (Logeswaran et al., 2019; Wu et al., 2020),  
100 we also presuppose that the mentions within the  
101 document have already been tagged.

102 **Data.** We use ESCO titles as ground truth labels,  
103 containing 13,890 skills.<sup>3</sup> Table 1 presents the  
104 train, dev, and test data in our experiments. We  
105 leverage the train set introduced by Decorte et al.  
106 (2023)<sup>4</sup> along with the dev and test sets provided  
107 in Decorte et al. (2022).<sup>5</sup> The train set is syn-  
108 thetically generated by Decorte et al. (2023) with  
109 the gpt-3.5-turbo-0301 model (OpenAI,  
110 2023). Specifically, this involves taking each skill  
111 from ESCO and prompting the model to generate  
112 sentences resembling JD sentences that require that  
113 particular skill. The dev and test splits, conversely,  
114 are derived from actual job advertisements sourced  
115 from the study by Zhang et al. (2022a). These  
116 JDs are annotated with spans corresponding to

<sup>3</sup>Per version 1.1.1, accessed on 01 August 2023.

<sup>4</sup><https://t.ly/edqkp>

<sup>5</sup><https://t.ly/LcqQ7>

117 specific skills, and these spans have subsequently  
118 been manually linked to ESCO, as described  
119 in the work of Decorte et al. (2022). In cases  
120 where skills cannot be linked, two labels are used,  
121 namely UNDESPECIFIED and LABEL NOT  
122 PRESENT. For the sake of uniformity, we map  
123 both of these labels to a generic UNK tag. We used  
124 several heuristics based on Levenshtein distance  
125 and sentence similarity to find the exact subspans  
126 if it exceeds certain thresholds, otherwise, it is  
127 UNK. This process is outlined in Appendix A.  
128 In addition, some data examples can be found  
129 in Appendix B. The number of UNKs in the data  
130 is also in Table 1. During inference, the UNK title  
131 is a prediction option for the models.

132 **Models.** We use two EL models, selected for  
133 their robust performance in EL on Wikipedia.<sup>6</sup>

134 **BLINK (Wu et al., 2020).** BLINK uses a bi-  
135 encoder architecture based on BERT (Devlin et al.,  
136 2019), for modeling pairs of mentions and entities.  
137 The model processes two inputs:

138 [CLS] ctxt<sub>l</sub> [S] mention [E] ctxt<sub>r</sub> [SEP]

139 Where “mention”, “ctxt<sub>l</sub>”, and “ctxt<sub>r</sub>” corresponds  
140 to the wordpiece tokens of the mention, the left  
141 context, and the right context. The mention is de-  
142 noted by special tokens [S] and [E]. The entity  
143 and its description are structured as follows:

144 [CLS] title [ENT] description [SEP]

145 Here, “title” and “description” represent the word-  
146 piece tokens of the skills’s title and description,  
147 respectively. [ENT] is a special token to separate  
148 the two representations. We train the model to max-  
149 imize the dot product of the [CLS] representation  
150 of the two inputs, for the correct skill in comparison  
151 to skills within the same batch. For each training  
152 pair  $(m_i, e_i)$ , the loss is computed as  $\mathcal{L}(m_i, e_i) =$   
153  $-s(m_i, e_i) + \log \sum_{j=1}^B \exp(s(m_i, e_j))$ , where  
154 the objective is to minimize the distance between  
155  $m_i$  and  $e_i$  while encouraging the model to assign a  
156 higher score to the correct pair and lower scores to  
157 randomly sampled incorrect pairs. Hard negatives  
158 are also used during training, these are obtained by  
159 finding the top 10 predicted skills for each training  
160 example. These extra hard negatives are added to  
161 the random in-batch negatives.

<sup>6</sup>For the hyperparameter setups, we refer to Appendix C.

	Train Source	Acc@1	Acc@4	Acc@8	Acc@16	Acc@32
Random		0.22±0.00	0.88±0.00	1.76±0.00	3.52±0.00	7.04±0.00
TF-IDF		2.25±0.00				
BLINK (bert-base)	ESCO	12.74±0.49	22.81±0.79	27.70±0.82	32.44±1.33	36.46±1.07
BLINK (bert-large)	ESCO	12.77±0.94	22.58±1.47	27.24±1.23	31.75±0.89	36.10±1.28
BLINK (bert-large)	Wiki (0-shot)	23.30±0.00	<b>32.89±0.00</b>	<b>38.16±0.00</b>	42.60±0.00	45.56±0.00
BLINK (bert-large)	Wiki + ESCO	<b>23.55±0.14</b>	32.63±0.16	37.38±0.09	<b>43.25±0.13</b>	<b>48.98±0.21</b>
GENRE (bart-base)	ESCO	1.47±0.05	4.84±1.74	10.46±6.81	11.30±4.18	15.51±4.62
GENRE (bart-large)	ESCO	2.33±0.44	5.74±1.43	8.18±2.21	11.13±2.42	15.26±2.66
GENRE (bart-large)	Wiki (0-shot)	6.91±0.00	12.34±0.00	15.52±0.00	21.60±0.00	33.17±0.00
GENRE (bart-large)	Wiki + ESCO	<b>11.48±0.41</b>	<b>21.26±0.43</b>	<b>27.40±0.78</b>	<b>37.21±0.69</b>	<b>49.78±1.05</b>

Table 2: **Skill Linking Results.** We show the results of the various models used. There are two base and four large models. Training sources are either ESCO or a combination of Wikipedia and ESCO. The results are the average and standard deviation over five seeds. For the 0-shot setup, we apply the fine-tuned models from the work of Wu et al. (2020) and Cao et al. (2021) to the ESCO test set once. We have a random and a TF-IDF-based baseline.

**GENRE (Cao et al., 2021).** GENRE formulates EL as a retrieval problem using a sequence-to-sequence model based on BART (Lewis et al., 2020). This model generates textual entity identifiers (i.e., skill titles) and ranks each entity  $e \in \mathcal{E}$  using an autoregressive approach:  $s(e | x) = p_\theta(y | x) = \prod_{i=1}^N p_\theta(y_i | y_{<i}, x)$ , where  $y$  represents the set of  $N$  tokens in the identifier of entity  $e$  (i.e., entity tile), and  $\theta$  denotes the model parameters. During decoding, the model uses a constrained beam search to ensure the generation of valid identifiers (i.e., only producing valid titles that exist within the KB, including UNK).

**Setup.** We train a total of six models: for BLINK, these are BERT<sub>base</sub> and BERT<sub>large</sub> (uncased; Devlin et al., 2019) trained on ESCO, and another large version trained on Wikipedia and ESCO sequentially. GENRE has the same setup, but then with BART (Lewis et al., 2020). Additionally, we apply the released models from both BLINK and GENRE (large, trained on Wikipedia) in a zero-shot manner and evaluate their performance. The reason we use Wikipedia-based models is that we hypothesize this is due to many skills in ESCO also having corresponding Wikipedia pages (e.g., Python<sup>7</sup> or teamwork<sup>8</sup>), thus could potentially help linking. Next, to address unknown entities (UNK), we include them as possible label outputs. Last, our evaluation metric is Accuracy@ $k$ , following prior research (Logeswaran et al., 2019; Wu et al., 2020; Zaporozhets et al., 2022).

<sup>7</sup>[https://en.wikipedia.org/wiki/Python\\_\(programming\\_language\)](https://en.wikipedia.org/wiki/Python_(programming_language))

<sup>8</sup><https://en.wikipedia.org/wiki/Teamwork>

### 3 Results

Table 2 presents the results. Each model is trained for five seeds, and we report the average with standard deviation. We make use of a random and TF-IDF-based baseline. Firstly, we observe that the strict linking performance (i.e., Acc@1) is rather modest for both BLINK and GENRE. But most models outperform the baselines. Notably, the top-performing models in this context are the BERT<sub>large</sub> and BART<sub>large</sub> models, which were further fine-tuned from Wikipedia EL with ESCO. As expected, scores improve considerably as we increase the value of  $k$ . Secondly, for both BLINK and GENRE, model size seems not to have a substantial impact when trained only on ESCO. Specifically for BLINK, the performance remains consistent for Acc@1 and exhibits only a slight decline as we relax the number of candidates for performance evaluation. For GENRE, the observed trend remains largely unchanged, even with a larger  $k$ .

Remarkably, the zero-shot setup performance of both BLINK and GENRE, when trained on Wikipedia, surpasses that of models trained solely on ESCO. For Wikipedia-based evaluation, GENRE usually outperforms BLINK. We notice the opposite in this case. For BLINK, this improvement is approximately 11 accuracy points for  $k = 1$ . Meanwhile, for GENRE, we observe an increase of roughly 9 accuracy points when trained on both Wikipedia and ESCO. This trend persists for a larger  $k$ , reaching up to a 12.5 accuracy point improvement for BLINK and a 34 accuracy point improvement for GENRE in the case of Acc@32. Furthermore, we show that further fine-tuning the Wikipedia-trained models on ESCO contributes to

Mention	BLINK	GENRE
① Work in a way that is <b>patient-centred</b> and inclusive.	person centred care (K0913)	work in an organised manner (T)
② You can <b>ride a bike</b> .	sell bicycles (S1.6.1)	drive two-wheeled vehicles (S8.2.2)
③ It is expected that you are a super user of the <b>MS office tools</b> .	use Microsoft Office (S5.6.1)	tools for software configuration management (0613)
④ <b>Picking and packing</b> .	carry out specialised packing for customers (S6.1.3)	perform loading and unloading operations (S6.2.1)
⑤ You are expected to be able to further <b>develop your team</b> - both personally and professionally. <b>GOLD: manage a team (S4.8.1)</b>	manage personal professional development (S1.14.1)	shape organisational teams based on competencies (S4.6.0)
⑥ Our games are developed using Unity so we expect all our programmers to have solid knowledge of mobile game development in Unity3D and <b>C#</b> .	C# (K0613)	C# (K0613)

Table 3: We show six qualitative examples. The mention is indicated with purple and we show the predictions ( $k = 1$ ) of BLINK and GENRE. Green predictions mean correct, and red indicates wrong linking with respect to the ground truth. We also show the ESCO ID, indicating the differences in concepts. The results show successful linking of implicit mentions of skills. In example (5), we show how the linked results are still valid while being different concepts. However, evaluation does not count it as a correct hit.

an improved EL performance at  $k = \{1, 16, 32\}$  for both models. For UNK-specific results, we refer to [Appendix D](#). We confirm our hypothesis that Wikipedia has concepts that are also in ESCO, this gives the model strong prior knowledge of skills.

## 4 Discussion

**Qualitative Analysis.** We manually inspected a subset of the predictions. We present qualitative examples in [Table 3](#). We found the following trends upon inspection:

- The EL models exhibit success in linking implicit and explicit mentions to their respective taxonomy titles (e.g., ①, ②, ④, ⑥).
- In cases of hard skills (③, ⑥), BLINK correctly matches “MS office tools” to “using Microsoft Office”, which is not an exact match. Both models predict the explicit mention “C#” correctly to the C# taxonomy title.
- We found that the models predict paraphrased versions of skills that could also be considered correct (④, ⑤), even being entirely different concepts (i.e., different ESCO IDs).

**Evaluation Limitation.** We qualitatively demonstrate the linking of skills that are implicit and/or valid. Empirically, we observe that the strict linking of skills leads to an underestimation of model performance. We believe this limitation is rooted in evaluation. In train, dev, and test, there is only

one correct gold label. We reciprocate the findings by [Li et al. \(2020\)](#), where they found that a large number of predictions are “technically correct” but limitations in Wikipedia-based evaluation falsely penalized their model (i.e., a more or less precise version of the same entity). Especially ⑤ in [Table 3](#) shows this challenge for ESCO, we can consider multiple links to be correct for a mention given a particular context. This highlights the need for appropriate EL evaluation sets, not only for ESCO, but for EL in general.

## 5 Conclusion

We present entity linking in the job market domain, using two existing high-performing neural models. We demonstrate that the bi-encoder architecture of BLINK is more suited to the job market domain compared to the autoregressive GENRE model. While strict linking results favor BLINK over GENRE, if we relax the number of candidates, we observe that GENRE performs slightly better. From a qualitative perspective, the performance of strict linking results is modest due to limitations in the evaluation set, which considers only one skill correct per mention. However, upon examining the predictions, we identify valid links, suggesting the possibility of multiple correct links for a particular mention, highlighting the need for more comprehensive evaluation. We hope this work sparks interest in entity linking within the job market domain.

## 285 Limitations

286 In the context of EL for ESCO, our approach has  
287 several limitations. Firstly, it only supports English,  
288 and might not generalize to other languages.

289 Secondly, our EL model is trained on synthetic  
290 training data, which may not fully capture the in-  
291 tricacies and variations present in real-world doc-  
292 uments. The use of synthetic data could limit its  
293 performance on actual, real JD texts. Nevertheless,  
294 we have human-annotated evaluation data.

295 Moreover, in our evaluation process, we use only  
296 one gold-standard ESCO title as the correct answer.  
297 This approach may not adequately represent a real-  
298 world scenario, where multiple ESCO titles could  
299 be correct as shown in Table 3.

## 300 References

301 Tom Ayoola, Shubhi Tyagi, Joseph Fisher, Christos  
302 Christodoulopoulos, and Andrea Pierleoni. 2022. [Re-  
303 FinED: An efficient zero-shot-capable approach to  
304 end-to-end entity linking](#). In *Proceedings of the 2022  
305 Conference of the North American Chapter of the  
306 Association for Computational Linguistics: Human  
307 Language Technologies: Industry Track*, pages 209–  
308 220, Hybrid: Seattle, Washington + Online. Associa-  
309 tion for Computational Linguistics.

310 Krisztian Balog, Yi Fang, Maarten De Rijke, Pavel  
311 Serdyukov, and Luo Si. 2012. [Expertise retrieval](#).  
312 *Foundations and Trends in Information Retrieval*,  
313 6(2–3):127–256.

314 Erik Brynjolfsson and Andrew McAfee. 2011. *Race  
315 against the machine: How the digital revolution is  
316 accelerating innovation, driving productivity, and  
317 irreversibly transforming employment and the econ-  
318 omy*. Brynjolfsson and McAfee.

319 Erik Brynjolfsson and Andrew McAfee. 2014. *The  
320 second machine age: Work, progress, and prosperity  
321 in a time of brilliant technologies*. WW Norton &  
322 Company.

323 de Nicola Cao, Gautier Izacard, Sebastian Riedel, and  
324 Fabio Petroni. 2021. [Autoregressive entity retrieval](#).  
325 In *International Conference on Learning Representa-  
326 tions*.

327 Nicola de Cao, Ledell Wu, Kashyap Popat, Mikel  
328 Artetxe, Naman Goyal, Mikhail Plekhanov, Luke  
329 Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and  
330 Fabio Petroni. 2022. Multilingual autoregressive en-  
331 tity linking. *Transactions of the Association for Com-  
332 putational Linguistics*, 10:274–290.

333 Benjamin Clavié and Guillaume Soulié. 2023. Large  
334 language models as batteries-included zero-shot esco  
335 skills matchers. *arXiv preprint arXiv:2307.03539*.

Jens-Joris Decorte, Jeroen Van Haute, Johannes  
Deleu, Chris Develder, and Thomas Demeester. 2022. Design of negative sampling strategies for  
distantly supervised skill extraction. *arXiv preprint  
arXiv:2209.05987*.

Jens-Joris Decorte, Severine Verlinden, Jeroen  
Van Haute, Johannes Deleu, Chris Develder, and  
Thomas Demeester. 2023. Extreme multi-label skill  
extraction training using large language models.  
*arXiv preprint arXiv:2307.10778*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and  
Kristina Toutanova. 2019. [BERT: Pre-training of  
deep bidirectional transformers for language under-  
standing](#). In *Proceedings of the 2019 Conference of  
the North American Chapter of the Association for  
Computational Linguistics: Human Language Tech-  
nologies, Volume 1 (Long and Short Papers)*, pages  
4171–4186, Minneapolis, Minnesota. Association for  
Computational Linguistics.

Octavian-Eugen Ganea and Thomas Hofmann. 2017. [Deep joint entity disambiguation with local neural  
attention](#). In *Proceedings of the 2017 Conference on  
Empirical Methods in Natural Language Processing*,  
pages 2619–2629, Copenhagen, Denmark. Associa-  
tion for Computational Linguistics.

Ann-sophie Gnehm, Eva Bühlmann, Helen Buchs, and  
Simon Clematide. 2022. [Fine-grained extraction  
and classification of skill requirements in German-  
speaking job ads](#). In *Proceedings of the Fifth Work-  
shop on Natural Language Processing and Computa-  
tional Social Science (NLP+CSS)*, pages 14–24, Abu  
Dhabi, UAE. Association for Computational Linguis-  
tics.

Zhengyan He, Shujie Liu, Mu Li, Ming Zhou, Longkai  
Zhang, and Houfeng Wang. 2013. [Learning entity  
representation for entity disambiguation](#). In *Proceed-  
ings of the 51st Annual Meeting of the Association  
for Computational Linguistics (Volume 2: Short Pa-  
pers)*, pages 30–34, Sofia, Bulgaria. Association for  
Computational Linguistics.

Martin le Vrang, Agis Papantoniou, Erika Pauwels,  
Pieter Fannes, Dominique Vandenstein, and Johan  
De Smedt. 2014. Escos: Boosting job matching in  
europe with semantic interoperability. *Computer*,  
47(10):57–64.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan  
Ghazvininejad, Abdelrahman Mohamed, Omer Levy,  
Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training  
for natural language generation, translation, and com-  
prehension](#). In *Proceedings of the 58th Annual Meet-  
ing of the Association for Computational Linguistics*,  
pages 7871–7880, Online. Association for Computa-  
tional Linguistics.

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 *Proceed-  
ings of the 2020 Conference on Empirical Methods*  
393

394					
395					
396					
397	Lajanugen Logeswaran, Ming-Wei Chang, Kenton Lee,				
398	Kristina Toutanova, Jacob Devlin, and Honglak Lee.				
399	2019. <a href="#">Zero-shot entity linking by reading entity de-</a>				
400	<a href="#">scriptions</a> . In <i>Proceedings of the 57th Annual Meet-</i>				
401	<i>ing of the Association for Computational Linguistics</i> ,				
402	pages 3449–3460, Florence, Italy. Association for				
403	Computational Linguistics.				
404	Yasumasa Onoe and Greg Durrett. 2020. Fine-grained				
405	entity typing for domain independent entity linking.				
406	In <i>Proceedings of the AAAI Conference on Artificial</i>				
407	<i>Intelligence</i> , volume 34, pages 8576–8583.				
408	OpenAI. 2023. <a href="#">Chatgpt (march version)</a> .				
409	Jonathan Raiman and Olivier Raiman. 2018. Deeptype:				
410	multilingual entity linking by neural type system evo-				
411	lution. In <i>Proceedings of the AAAI Conference on</i>				
412	<i>Artificial Intelligence</i> , volume 32.				
413	Nils Reimers and Iryna Gurevych. 2019. <a href="#">Sentence-</a>				
414	<a href="#">BERT: Sentence embeddings using Siamese BERT-</a>				
415	<a href="#">networks</a> . In <i>Proceedings of the 2019 Conference on</i>				
416	<i>Empirical Methods in Natural Language Processing</i>				
417	<i>and the 9th International Joint Conference on Natu-</i>				
418	<i>ral Language Processing (EMNLP-IJCNLP)</i> , pages				
419	3982–3992, Hong Kong, China. Association for Com-				
420	putational Linguistics.				
421	Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian				
422	Riedel, and Luke Zettlemoyer. 2020. <a href="#">Scalable zero-</a>				
423	<a href="#">shot entity linking with dense entity retrieval</a> . In				
424	<i>Proceedings of the 2020 Conference on Empirical</i>				
425	<i>Methods in Natural Language Processing (EMNLP)</i> ,				
426	pages 6397–6407, Online. Association for Computa-				
427	tional Linguistics.				
428	Klim Zaporozhets, Lucie-Aimée Kaffee, Johannes Deleu,				
429	Thomas Demeester, Chris Develder, and Isabelle Au-				
430	genstein. 2022. Tempel: Linking dynamically evolv-				
431	ing and newly emerging entities. <i>Advances in Neural</i>				
432	<i>Information Processing Systems</i> , 35:1850–1866.				
433	Mike Zhang, Kristian Jensen, Sif Sonniks, and Barbara				
434	Plank. 2022a. <a href="#">SkillSpan: Hard and soft skill extrac-</a>				
435	<a href="#">tion from English job postings</a> . In <i>Proceedings of</i>				
436	<i>the 2022 Conference of the North American Chap-</i>				
437	<i>ter of the Association for Computational Linguistics:</i>				
438	<i>Human Language Technologies</i> , pages 4962–4984,				
439	Seattle, United States. Association for Computational				
440	Linguistics.				
441	Mike Zhang, Kristian Nørgaard Jensen, and Barbara				
442	Plank. 2022b. <a href="#">Kompetencer: Fine-grained skill clas-</a>				
443	<a href="#">sification in Danish job postings via distant super-</a>				
444	<a href="#">vision and transfer learning</a> . In <i>Proceedings of the</i>				
445	<i>Thirteenth Language Resources and Evaluation Con-</i>				
446	<i>ference</i> , pages 436–447, Marseille, France. European				
447	Language Resources Association.				
	Fangwei Zhu, Jifan Yu, Hailong Jin, Lei Hou, Juanzi				
	Li, and Zhifang Sui. 2023. <a href="#">Learn to not link: Explor-</a>				
	<a href="#">ing NIL prediction in entity linking</a> . In <i>Findings of</i>				
	<i>the Association for Computational Linguistics: ACL</i>				
	2023, pages 10846–10860, Toronto, Canada. Associ-				
	ation for Computational Linguistics.				
	<b>A Data Preprocessing</b>				
	We outline the preprocessing steps for the training				
	set. In <a href="#">Decorte et al. (2023)</a> , there are sentence–				
	ESCO skill title pairs. The data is synthetically				
	generated by GPT-3.5. Where for each ESCO skill				
	title a set of 10 sentences is generated. A crucial				
	limitation for entity linkers is that the generated				
	sentence does not have the ESCO skill title as an				
	exact match in the sentence, but at most slightly				
	paraphrased. To find the most similar subspan in				
	the sentence to the target skill, we have to apply				
	some heuristics. In <a href="#">Algorithm 1</a> , we denote our				
	algorithm to find the most similar subspan. Our				
	method is a brute force approach, where we cre-				
	ate all possible n-grams until the maximum length				
	of the sentence, and compare the target subspan				
	against each n-gram. Based on Levenshtein dis-				
	tance, we filter the results, where we only take the				
	top 80% n-grams. Then, we encode both target				
	subspan and n-gram with SentenceBERT ( <a href="#">Reimers</a>				
	<a href="#">and Gurevych, 2019</a> ), the similarity is based on				
	cosine similarity. If the similarity does not exceed				
	0.5, the candidate subspan is UNK and the ESCO				
	title will also be UNK, otherwise, we take the most				
	similar n-gram. Empirically, we found that these				
	thresholds worked best. Note that this method is				
	not error-prone, but allows us to generate implicit				
	and negative examples to train entity linkers. We				
	show two qualitative examples in <a href="#">Figure 1</a> and dis-				
	cuss the quality in <a href="#">Appendix B</a> .				
	<b>B Data Examples</b>				
	We show a couple of data examples from the train-				
	ing ( <a href="#">Figure 1</a> ) and development set ( <a href="#">Figure 2</a> ). In				
	the training examples, we show an example with				
	a mention that is the same as the original ESCO				
	title (“young horse training”). In addition, we have				
	an example where there is an “implicit” mention				
	(i.e., the mention does not exactly match with the				
	label title). This shows that our algorithm works				
	to an extent. For the development example, this is				
	another implicit mention. However, these samples				
	are human annotated. There are also quite some				
	UNKs given the training data. We show that this is				
	helping the model predict UNK.				

---

**Algorithm 1:** Find the most similar n-gram to a target subspan

---

**Data:** *sentence*: The input sentence

*target\_subspan*: The target subspan

*threshold*: The Levenshtein distance similarity threshold

**Result:** *most\_similar\_ngram*: The most similar n-gram

```
1 all_ngrams ← GenerateAllNgrams(sentence)
2 filtered_ngrams ← FilterNgrams(all_ngrams, target_subspan, threshold)
3 most_similar_ngram ← None
4 max_similarity ← 0
5 for ngram in filtered_ngrams do
6   subspan_embedding ← EncodeWithSBERT(target_subspan)
7   ngram_embedding ← EncodeWithSBERT(ngram)
8   similarity ← CosineSimilarity(subspan_embedding, ngram_embedding)
9   if similarity > max_similarity and similarity > 0.5 then
10    | max_similarity ← similarity
11    | most_similar_ngram ← ngram
12  else
13    | most_similar_ngram = UNK
14 return most_similar_ngram
```

---

Table 4: **UNK Linking Results.** We show the results of BLINK and GENRE predicting UNK. We use the best-performing models, based on Table 2.

	Train Source	Acc@1	Acc@4	Acc@8	Acc@16	Acc@32
BLINK (bert-large) UNK	Wiki + ESCO	1.38±0.12	3.32±0.22	4.67±0.33	7.68±0.42	10.70±0.58
GENRE (bart-large) UNK	Wiki + ESCO	1.65±0.20	4.99±0.50	9.23±0.58	16.01±0.48	24.70±2.52

## C Implementation Details

For training both BLINK<sup>9</sup> and GENRE,<sup>10</sup> we use their respective repositories. All models are trained for 10 epochs, for a batch size of 32 for training and 8 for evaluation. For both BLINK and GENRE we use 5% warmup. For the base models we use learning rate  $2 \times 10^{-5}$  and for the large models we use  $2 \times 10^{-6}$ . The maximum context and candidate length is 128 for both models. Each model is trained on an NVIDIA A100 GPU with 40GBs of VRAM and an AMD Epyc 7662 CPU. The seed numbers the models are initialized with are 276800, 381552, 497646, 624189, 884832. We run all models with the maximum number of epochs (10) and select the best-performing one based on validation set performance for accuracy@1.

<sup>9</sup><https://github.com/facebookresearch/BLINK>

<sup>10</sup><https://github.com/facebookresearch/genre>

## D UNK Evaluation

In Table 4, we show the performance of both BLINK and GENRE on the UNK label. We use the best-performing models based on Table 2. Generally, we observe that GENRE is better in predicting UNKs than BLINK. However, the exact linking results (i.e., Acc@1) are low. This can potentially be alleviated by actively training for predicting UNKs (Zhu et al., 2023).

```

1  {
2    "context_left": "we're looking for someone who is passionate
3    about",
4    "context_right": "and eager to share their knowledge with
5    others.",
6    "mention": "young horse training",
7    "label_title": "young horses training",
8    "label": "Principles & techniques of educating young horses
9    important simple body control exercises.",
10   "label_id": 2198
11  }
12  {
13   "context_left": "Hands-on experience with",
14   "context_right": "is a must-have qualification for this
15   job.",
16   "mention": "various hand-operated printing devices",
17   "label_title": "types of hand-operated printing devices",
18   "label": "Process of creating various types hand-operated
19   printing devices, such as stamps, seals, embossing labels or
20   inked pads and their applications.",
21   "label_id": 10972
22  }

```

Figure 1: **Two Training Examples.** The training examples are in the format for BLINK, there is the left context, right context, and the mention. The label title is the ESCO skill, and the label is the description of the label title. The label ID is the ID that refers to the label title.

```

1  {
2    "context_left": "You must have an",
3    "context_right": "with a high-quality mindset.",
4    "mention": "analytical proactive and structured workstyle",
5    "label_title": "work in an organised manner",
6    "label": "Stay focused on the project at hand, at any time.
7    Organise, manage time, plan, schedule and meet deadlines.",
8    "label_id": 3884
9  }

```

Figure 2: **One Evaluation Example.** The evaluation example is in the format for BLINK, there is the left context, right context, and the mention. The label title is the ESCO skill, and the label is the description of the label title. The label ID is the ID that refers to the label title.