

Learning Job Title Representation from Job Description Aggregation Network

Napat Laosaengpha[♣], Thanit Tativannarat[♣], Chawan Piansaddhayanon[♡]

Attapol Rutherford[◇], and Ekapol Chuangsuwanich^{♣, ♡}

[♣]Department of Computer Engineering, Faculty of Engineering, Chulalongkorn University

[♡]Center of Excellence in Computational Molecular Biology,

Faculty of Medicine, Chulalongkorn University

[◇]Department of Linguistics, Faculty of Arts, Chulalongkorn University, Thailand

napatnicky@gmail.com thanit.tati@gmail.com schwanph@gmail.com

attapol.t@chula.ac.th ekapolc@cp.eng.chula.ac.th

Abstract

Learning job title representation is a vital process for developing automatic human resource tools. To do so, existing methods primarily rely on learning the title representation through skills extracted from the job description, neglecting the rich and diverse content within. Thus, we propose an alternative framework for learning job titles through their respective job description (JD) and utilize a Job Description Aggregator component to handle the lengthy description and bidirectional contrastive loss to account for the bidirectional relationship between the job title and its description. We evaluated the performance of our method on both in-domain and out-of-domain settings, achieving a superior performance over the skill-based approach.

1 Introduction

With the rapid expansion of online recruitment platforms, vast amounts of job advertisement data (JAD) have been generated. One key part of this data is a job posting, providing detailed information on job titles, specialties, and responsibilities for open positions. Thus, the availability of a system that could understand the post semantics, especially job titles, would greatly facilitate the matchmaking process between both the recruiters and job applicants. This leads to a surge of interest in learning job title representation due to its potential ability to automate job-related tasks such as job recommendation (Kaya and Bogers, 2021; Zhao et al., 2021), job trajectory prediction (Decorte et al., 2023a), and job title benchmarking (Zhang et al., 2019).

To learn the title representation, previous works have primarily relied on utilizing skills information to learn the association between the job title and their respective skill (Decorte et al., 2021; Zbib et al., 2022; Bocharova et al., 2023). However, this approach also has some shortcomings as it requires skill information. The skills for a given job are

either manually listed, which can be erroneous or incomplete, or automatically extracted from the job description through methods such as keyword matching or automatic skill extraction (Zhang et al., 2022b; Li et al., 2023). These skill extraction methods often require a predefined skill vocabulary or a curated dataset (Zhang et al., 2022a). Furthermore, it is necessary to keep these resources up-to-date with trends in the job market as the dynamic and rapid growth of emerging job roles.

Previous works mitigate these problems by generating synthetic skill data (Decorte et al., 2023b; Clavié and Soulié, 2023) or creating datasets where both job titles, and skill lists are readily available (Bhola et al., 2020; Goyal et al., 2023). Nonetheless, the former approach further increases pipeline complexity, while the latter suffers from missing skills annotation caused by a communication gap between employers and recruiters.

In this work, we propose to overcome the challenges of obtaining a comprehensive set of skills by bypassing the whole process and instead develop a new framework to learn job titles directly through job descriptions (JDs) without the need for the skill extraction pipeline. We introduce job description aggregation network, which reweights each segment of the JD by their importance and then aggregates them into a unified JD representation. Our contributions are as follows:

- Our JD-based method outperforms all previous skill-based approaches in both in-domain and out-of-domain settings, achieving up to a 1.8% and 1.0% absolute performance gain.
- Our ablation study shows that the ability to reweight segments according to their importance afforded by our model is critical to its accuracy.
- We show that our approach can implicitly learn the information about the underlying skills associated with job titles.

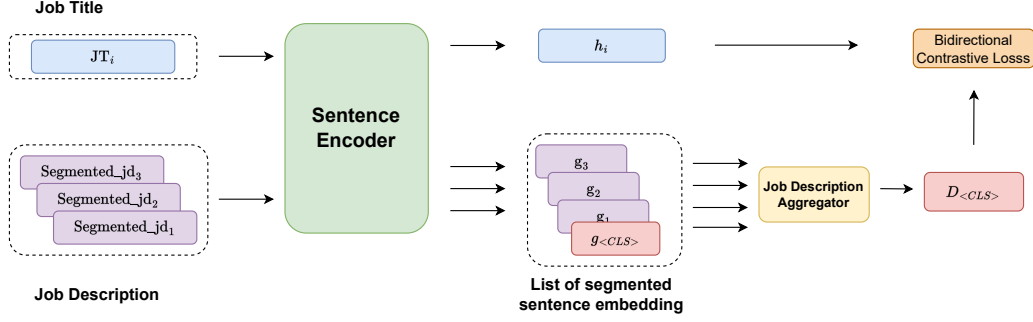


Figure 1: The overall proposed method, JD Aggregation Network (JDAN), presents a dual-encoder architecture coupled with a job description aggregator module.

2 Our Proposed Method

An overview of the proposed framework is illustrated in Figure 1. Initially, job titles and their respective segmented job descriptions are independently fed into a sentence encoder to obtain their representation. Subsequently, the sentence embeddings are fused through the job description aggregator to acquire a unified representation. Finally, bidirectional contrastive loss is utilized as a training objective to maximize a pairwise similarity between the embedded job title and their respective aggregated job description representation while minimizing others.

The subsections describe the sentence encoder, job description aggregator, and contrastive learning process in detail.

2.1 Sentence Encoder

The sentence encoder follows a dual-encoder architecture (Guo et al., 2018; Feng et al., 2022), in which the job title and description are fed into the encoder separately to generate their respective representations. The job title is fed into the model to obtain a job title embedding h . On the other hand, the job description can be lengthy with some irrelevant parts. Thus, the job description is broken down into sentences and encoded sentence by sentence, resulting in a list of segmented sentence embeddings $G = [g_1, g_2, \dots, g_n]$. Then, the embeddings are aggregated into a final representation for the description. The sentence segmentation process is explained in Appendix A.

2.2 Job Description Aggregation Network (JDAN)

The job description aggregator is responsible for combining multiple sentence embeddings into a unified representation by weighting the importance

of each sentence. This step is important because job descriptions often contain information not directly related to the corresponding job titles such as location and salary. Inspired by (Wu et al., 2021; Zhang et al., 2023), we create an additional learnable token $g_{<CLS>}$ to represent the summarized token and prepend it to the sequence of sentence embedding $G = [g_{<CLS>}, g_1, g_2, \dots, g_n]$. A Layer Normalization and a shallow transformer encoder are then applied to D to obtain a list of learned representation $D = [d_{<CLS>}, d_1, d_2, \dots, d_n]$.

$$D = \text{TransformerEncoder}(\text{LayerNorm}(G)) \quad (1)$$

The learned summarized token $d_{<CLS>}$ is then fed through three MLP layers with ReLU activation to obtain a final unified representation f .

2.3 Bidirectional Contrastive Learning

The bidirectional contrastive learning minimizes the following training objective function:

$$\mathcal{L}_i = -\left(\log \frac{e^{\text{sim}(h_i, f_i)/\tau}}{\sum_{j=1}^N e^{\text{sim}(h_i, f_j)/\tau}} + \log \frac{e^{\text{sim}(h_i, f_i)/\tau}}{\sum_{j=1}^N e^{\text{sim}(h_j, f_i)/\tau}}\right) \quad (2)$$

where h_i is the i^{th} job title embedding, f_i is i^{th} the unified job description embedding, τ is the temperature scaling parameter, and N is the batch size. Instead of only maximizing the similarity between the job title embedding h_i and its respective job description embedding f_i while minimizing the similarities of other pairs $\text{sim}(h_i, f_j)$, the objective function also further introduces the second term to minimize the similarity of the job description embedding f_i to other job titles $\text{sim}(h_j, f_i)$. This could be seen as an extension of SimCSE (Gao et al., 2021) where an additional latter term is included to mitigate the overlooked characteristic of

JAD where a bidirectional relationship (job title to job description and vice versa) exists. (Yang et al., 2019; Feng et al., 2022).

3 Experimental Setup

We benchmarked the performance of our proposed framework under two settings: **in-domain** and **out-of-domain**.

For the in-domain evaluation, we used our own **JTG-Jobposting** dataset for training and **JTG-Synonym** for evaluation. JTG-Jobposting is a private Thai-English job posting dataset consisting of 28,844 job postings from, <https://jobtopgun.com>, a renowned recruitment website in Thailand. The postings include job titles, job descriptions, and skills. We performed benchmarking on the JTG-Synonym dataset by posing the problem as a cross-lingual synonym retrieval task where the job titles were used as queries, and all synonyms were used as the candidate pool. Each query was performed on the English and Thai candidate pools separately to calculate the R@5, R@10, and mAP@25. The final metric values were obtained by averaging across every query-candidate-pool pairs. This evaluation protocol was intentionally designed to avoid language bias where a query would prefer a candidate from the same language. This issue will be later discussed in the Section 4.5.

For the out-of-domain evaluation, we used **My-careersfuture.sg** (Bhola et al., 2020), a dataset of real-world job postings collected by the Singaporean government as the training set. ESCO (le Vrang et al., 2014), a standardized system of the European Union (EU) for categorizing skills, competencies, qualifications and occupations was used as the validation and test data. The task chosen for benchmarking is job normalization where the goal is to predict the standardized version of each job title. The ESCO job normalization dataset consists of 30,926 unique job titles and 2,675 standardized ESCO occupation labels. We followed (Decorte et al., 2021) and used the standard micro-average of recall at 5, 10 and MRR as metrics.

The summary statistics of the datasets are shown in Table 3. See Appendix C for more details.

3.1 Implementation Details

We compared the performance of our proposed framework to other previously proposed skill-based methods, which are **JobBERT** (Decorte et al.,

2021), **Doc2VecSkill** (Zbib et al., 2022), **VacancySBERT** (Bocharova et al., 2023), and our own skill-based approach. In our approach, we used a dual-encoder to independently encode the job title and the concatenated set of skills corresponding to the title using a specified separator token ("[SEP]" for BERT (Devlin et al., 2019) and "</s>" for XLM-R (Conneau et al., 2019)). Then, the bidirectional contrastive loss was used to encourage the job title and the respective encoded skills pair to come together while pushing the others away. All skill-based models used keyword matching to extract skills from the job description.

For automatic skill extraction, we used SkillSpan¹ to extract keywords from the Mycareersfuture.sg dataset. On the other hand, as the JTG-job posting dataset contains both English and Thai, we made some modifications by creating a new classifier to extract the skills from the job description by posing the problem as a multi-label classification instead. Additional details for the automatic skill extraction is provided in Appendix D.2.

The experiments were conducted using the pretrained-language model BERT as a sentence encoder (Devlin et al., 2019) on the job normalization task and XLM-R (Conneau et al., 2019) on the synonym retrieval task because JTG-Jobposting contains both Thai and English. The sentence representations of every model were obtained through mean pooling of every token (words) in the sentence. We reported the average with standard deviation from five random seeds. The hyperparameters were obtained through grid search in the validation set. The extended implementation detail of our method is shown in the Appendix B, D.

4 Results and Analysis

4.1 Main Results

Our method achieved consistent performance improvement over all previous skill-based approaches on both datasets under all metrics, achieving up to 1.8% and 1.0% absolute performance gain on the JTG-Synonym and Job normalization task, respectively (Table 1).

¹https://huggingface.co/jjzha/jobbert_skill_extraction

Method	JTG-Synonym			Job Normalization		
	R@5 ↑	R@10 ↑	mAP@25 ↑	R@5 ↑	R@10 ↑	MRR ↑
XLM-R / BERT (w/o finetuning)	15.11	18.68	10.27	26.23	32.10	20.57
Skill-based method						
JobBERT (Decorte et al., 2021)	31.04 (±0.61)	41.99 (±0.86)	22.83 (±0.42)	38.65 [†]	46.04 [†]	30.92 [†]
Doc2VecSkill (Zbib et al., 2022)	25.64 (±0.35)	34.44 (±0.61)	18.50 (±0.28)	45.95 [†]	54.00 [†]	34.14 [†]
VacancySBERT (Bocharova et al., 2023)	-	-	-	42.50 [†]	55.60 [†]	-
Keyword Skill (ours)	<u>47.79</u> (±0.32)	<u>60.62</u> (±0.32)	<u>35.87</u> (±0.13)	<u>48.23</u> (±0.19)	<u>56.45</u> (±0.18)	<u>37.68</u> (±0.15)
JD-based method						
JD Aggregation Network (ours)	49.21 * (±0.55)	62.34 * (±0.31)	37.09 * (±0.45)	49.24 * (±0.43)	57.22 * (±0.43)	38.71 * (±0.24)

Table 1: The performance of our proposed JD-based method against other skill-based methods. The best results are bolded, and the second-best ones are underlined. †: results from their original papers. * denotes significant improvement over Keyword Skill (ours) using two-sample t-test ($p < 0.01$).

4.2 Comparison between Skill-based and JD-based method

We found that our framework also outperformed or achieved competitive performance when compared to other skill-based approaches, even with human annotations (Table 2). Surprisingly, our method performed better in skill-based recruiter annotation on the job normalization task but not the JTG-Synonym task. This is because skill annotations for the MyCareersFuture.sg dataset might be incomplete due to a communication gap between employers and recruiters (Bhola et al., 2020). In addition, some parts of the JTG-Jobposting dataset also contain implicit information that is only present in the recruiter’s annotated skill and not explicitly mentioned in the JD. For example, a job posting for "Data analyst" is annotated with the skills "Excel", "SQL", and "Python," but these skills do not appear in any part of the JD. Additional examples could be seen in Figure 3 of our Appendix.

	JTG-Synonym	Job Normalization
	R@10 ↑	R@10 ↑
Skill-based		
- Recruiter Annotation	63.06 (±0.67)	51.51 (±0.30)
- Keyword Matching	60.62 (±0.32)	<u>56.45</u> (±0.18)
- Skill Extraction	51.84 (±0.87)	55.11 (±0.27)
JD-based (ours)	<u>62.34</u> (±0.31)	57.22 (±0.43)

Table 2: Comparison between learning job titles through skills or job descriptions.

4.3 Probe Analysis on Job Title Embeddings

We also conducted an additional analysis on the learned embeddings by linear probing (Alain and Bengio, 2017) (training a linear classifier on top of the embeddings). We trained classifiers to predict skills from the job title embeddings learned from various methods. As shown in Figure 2, the

result of linear probing on a held-out set of the JTG-Jobposting dataset (details provided in Appendix D.1) found that the model learned through recruiter-annotated skills and JD were on par (25.8 vs 24.5 Top-10 accuracy) which doubled the performance of using XLM-R embedding without finetuning (13.1 Top-10 accuracy), implying that the JD-based method could implicitly understand skill information despite not being trained with one. This offers an explanation on why learning the job title representation from the job description can be more beneficial than from skills as the skill information can be learned implicitly while having access to other additional information.

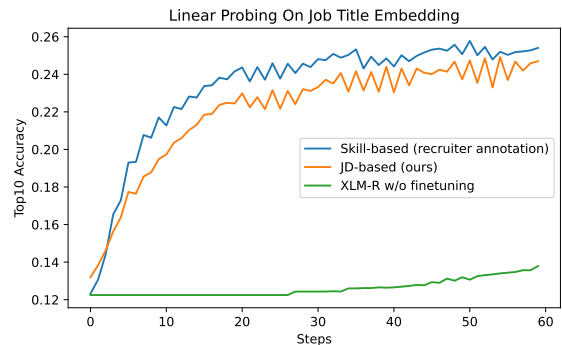


Figure 2: Comparison of job title linear probing from the model learning from the skill-based method, our JD-based method, and without any fine-tuning.

4.4 Aggregation Design Choices

Next, we analyze our design of the job description aggregator and compare it against three other possibilities (Table 5). Instead of using transformers to aggregate multiple sentence embeddings, we can use mean or max pooling. Another possibility is to use the encoder to encode the entire job description (Document Level). Our proposed method outperforms the others, highlighting the importance of having a segmented sentence representation and

	JTG-Jobposting	Mycarreersfuture.sg
# of job postings	28,844	20,298
# of total skills	301,124	405,606
# of distinct skill	35,107	2,548
# of skills with more than 20 occurrences	3,092	1,209
Average skill tags per job posting	10.43 \pm 19.91	19.98 \pm 0.06
Average token count of job description per job posting	123.10 \pm 82.13	155.77 \pm 103.55
Average # of segmented sentence per job posting	5.88 \pm 3.87	6.63 \pm 27.14
Average token count of segmented sentence	21.23 \pm 26.58	27.19 \pm 27.14

Table 3: Dataset statistics of JTG-Jobposting and Mycarreersfuture.sg.

Candidate pool	Thai pool			English pool			Combined pool		
	EN	TH	CS	EN	TH	CS	EN	TH	CS
XLM-R	2.56	34.84	29.58	31.46	3.05	8.24	17.00	19.73	15.35
JobBERT	29.29	57.11	50.85	49.59	27.72	48.53	30.78	32.44	33.44
Skill-based (ours)	56.14	68.05	62.35	64.22	52.59	59.13	35.35	41.57	38.24
JD-based (ours)	59.87	71.15	72.35	64.12	56.46	70.40	37.93	42.80	40.92

Table 4: The comparison of cross-lingual performance ($R@10 \uparrow$) on JTG-Synonym retrieval task. "CS" refers to queries which contain Thai and English code-switching.

weighting mechanism for each sentence. We provide examples of how the first attention layer selects important parts of the job description in Figure 6 and Section E of our Appendix.

	JTG-Synonym	Job Normalization
	$R@10 \uparrow$	$R@10 \uparrow$
Document Level	61.44 (± 0.51)	57.09 (± 0.39)
Sentence Level		
- max pooling	62.18 (± 0.30)	55.94 (± 0.29)
- mean pooling	61.42 (± 0.36)	56.66 (± 0.41)
- JD aggregator (ours)	62.34 (± 0.31)	57.22 (± 0.43)

Table 5: A Comparison of our proposed method (JD aggregator) against the other three baselines.

4.5 Cross-lingual Evaluation

Studies have pointed out the problem of language bias in textual embeddings (Roy et al., 2020; Yang et al., 2021). Embeddings from the same language are generally closer together compared to their cross-lingual counterparts. As a result, a query in Thai will prefer Thai candidates over English ones, and vice versa. To avoid this bias, the candidate pool was divided into Thai and English pool, and retrieval was done separately. Table 4 shows the retrieval results for different query-candidate-pool pairs. As expected, all models perform better despite the pool being split. Cross-lingual setups are generally more challenging. However, our method consistently outperforms others in every setting.

5 Conclusion

In this paper, we propose a framework for learning the semantic similarity of job titles through job descriptions, bypassing the need for a complete set of skills. The job description aggregator and bidirectional contrastive loss are also introduced to handle the nature of lengthy job descriptions and the two-way relationship between the job title and its description. Our results show that our approach achieves superior performance over the previous state-of-the-art skill-based methods.

6 Limitations

The limitations of our work are as follows:

- Our framework is limited to information only available in the job description. Thus, in some cases, our performance might be sub-optimal than recruiter annotations, which could provide information that is not explicitly mentioned in the job description.
- Our job description aggregator requires the description to be segmented. This could be challenging when applied to languages other than English.
- The job description aggregator is not designed for encoding the entire job description, which does not guarantee that our job description aggregator can be further used in downstream tasks that require job description embedding.

7 Ethics Statement

An inclusion of the job description for learning the job title could induce gender and age bias into job recommendations and search results. This might affect the fairness and inclusiveness of the job match-making process. (Saxena and Jain, 2021; Seyed-salehi et al., 2022).

8 Acknowledgements

This work is supported in part by JOBTOPGUN, job postings and recruitment platform in Thailand. We also would like to thank the Chulalongkorn Computational Molecular Biology Group (CMB@CU) for providing additional computational resources.

References

- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Akshay Bhola, Kishalay Halder, Animesh Prasad, and Min-Yen Kan. 2020. [Retrieving skills from job descriptions: A language model based extreme multi-label classification framework](#). In *International Conference on Computational Linguistics*.
- Maiia Y Bocharova, Eugene V Malakhov, and Vitaliy I Mezhuiev. 2023. [VacancySBERT: the approach for representation of titles and skills for semantic similarity search in the recruitment domain](#). *Applied Aspects of Information Technology*, 1(6):52–59.
- Benjamin Clavié and Guillaume Soulié. 2023. [Large language models as batteries-included zero-shot ESCO skills matchers](#). In *Proceedings of the 3rd Workshop on Recommender Systems for Human Resources (RecSys in HR 2023) co-located with the 17th ACM Conference on Recommender Systems (RecSys 2023), Singapore, Singapore, 18th-22nd September 2023*, volume 3490 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Jens-Joris Decorte, Jeroen Van Haute, Johannes Deleu, Chris Develder, and Thomas Demeester. 2023a. [Career path prediction using resume representation learning and skill-based matching](#). In *Proceedings of the 3rd Workshop on Recommender Systems for Human Resources (RecSys in HR 2023) co-located with the 17th ACM Conference on Recommender Systems (RecSys 2023), Singapore, Singapore, 18th-22nd September 2023*, volume 3490 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Jens-Joris Decorte, Jeroen Van Haute, Thomas Demeester, and Chris Develder. 2021. [JobBERT: Understanding job titles through skills](#). *CoRR*, abs/2109.09605.
- Jens-Joris Decorte, Severine Verlinden, Jeroen Van Haute, Johannes Deleu, Chris Develder, and Thomas Demeester. 2023b. [Extreme multi-label skill extraction training using large language models](#). *CoRR*, abs/2307.10778.
- 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-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nidhi Goyal, Jushaan Kalra, Charu Sharma, Raghava Mutharaju, Niharika Sachdeva, and Ponnurangam Kumaraguru. 2023. [JobXMLC: EXtreme multi-label classification of job skills with graph neural networks](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2181–2191, Dubrovnik, Croatia. Association for Computational Linguistics.
- Mandy Guo, Qinlan Shen, Yinfei Yang, Heming Ge, Daniel Cer, Gustavo Hernandez Abrego, Keith Stevens, Noah Constant, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Effective parallel corpus mining using bilingual sentence embeddings](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 165–176, Brussels, Belgium. Association for Computational Linguistics.
- Mesut Kaya and Toine Bogers. 2021. [Effectiveness of job title based embeddings on résumé to job ad recommendation](#). In *Proceedings of the Workshop on*

- Recommender Systems for Human Resources (RecSys in HR 2021) co-located with the 15th ACM Conference on Recommender Systems (RecSys 2021)*, Amsterdam, The Netherlands, 27th September - 1st October 2021, volume 2967 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Martin le Vrang, Agis Papantoniou, Erika Pauwels, Pieter Fannes, Dominique Vandenstein, and Johan De Smedt. 2014. [ESCO: Boosting job matching in europe with semantic interoperability](#). *Computer*, 47(10):57–64.
- Nan Li, Bo Kang, and Tijn De Bie. 2023. [SkillGPT: A RESTful API service for skill extraction and standardization using a large language model](#). *CoRR*, abs/2304.11060.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Uma Roy, Noah Constant, Rami Al-Rfou, Aditya Barua, Aaron Phillips, and Yinfei Yang. 2020. [LAReQA: Language-agnostic answer retrieval from a multilingual pool](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5919–5930, Online. Association for Computational Linguistics.
- Shrikant Saxena and Shweta Jain. 2021. [Exploring and mitigating gender bias in recommender systems with explicit feedback](#). *CoRR*, abs/2112.02530.
- Shirin Seyedsalehi, Amin Bigdeli, Negar Arabzadeh, Bhaskar Mitra, Morteza Zihayat, and Ebrahim Bagheri. 2022. [Bias-aware fair neural ranking for addressing stereotypical gender biases](#). In *Proceedings of the 25th International Conference on Extending Database Technology, EDBT 2022, Edinburgh, UK, March 29 - April 1, 2022*, pages 2:435–2:439. OpenProceedings.org.
- Christophe Van Gysel and Maarten de Rijke. 2018. [Pytreceval: An extremely fast python interface to trec_eval](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 873–876.
- Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. 2021. Representing long-range context for graph neural networks with global attention. *Advances in Neural Information Processing Systems*, 34:13266–13279.
- Yinfei Yang, Gustavo Hernández Ábrego, Steve Yuan, Mandy Guo, Qinlan Shen, Daniel Cer, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Improving multilingual sentence embedding using bi-directional dual encoder with additive margin softmax](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5370–5378. ijcai.org.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-hsuan Sung, Brian Strope, and Ray Kurzweil. 2020. [Multilingual universal sentence encoder for semantic retrieval](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 87–94, Online. Association for Computational Linguistics.
- Ziyi Yang, Yinfei Yang, Daniel Cer, Jax Law, and Eric Darve. 2021. [Universal sentence representation learning with conditional masked language model](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6216–6228, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rabih Zbib, Lucas Lacasa Alvarez, Federico Retyk, Rus Poves, Juan Aizpuru, Hermenegildo Fabregat, Vaidotas Simkus, and Emília Garcia Casademont. 2022. [Learning job titles similarity from noisy skill labels](#). *CoRR*, abs/2207.00494.
- Denghui Zhang, Junming Liu, Hengshu Zhu, Yanchi Liu, Lichen Wang, Pengyang Wang, and Hui Xiong. 2019. [Job2Vec: Job title benchmarking with collective multi-view representation learning](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2763–2771.
- Mike Zhang, Kristian Jensen, Sif Sonniks, and Barbara Plank. 2022a. [SkillSpan: Hard and soft skill extraction from English job postings](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4962–4984, Seattle, United States. Association for Computational Linguistics.
- Mike Zhang, Kristian Nørgaard Jensen, Rob van der Goot, and Barbara Plank. 2022b. [Skill extraction from job postings using weak supervision](#). In *Proceedings of the 2nd Workshop on Recommender Systems for Human Resources (RecSys-in-HR 2022) co-located with the 16th ACM Conference on Recommender Systems (RecSys 2022)*, Seattle, USA, 18th-23rd September 2022, volume 3218 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Shunyu Zhang, Yaobo Liang, Ming Gong, Daxin Jiang, and Nan Duan. 2023. [Modeling sequential sentence relation to improve cross-lingual dense retrieval](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Jing Zhao, Jingya Wang, Madhav Sigdel, Bopeng Zhang, Phuong Hoang, Mengshu Liu, and Mohammed Korayem. 2021. [Embedding-based recommender system for job to candidate matching on scale](#). *CoRR*, abs/2107.00221.

A Data Pre-processing

For the JTG-Jobposting dataset, we used the field "Job_Description" as a job description, and a heuristic algorithm was then applied to segment it by splitting them based on bullet points, hyphens, and numbering using regular expression.

For the MyCareersFuture.sg dataset, we concatenated the fields "Role & Responsibilities" and "Job Requirement" to represent the job description as suggested in their work. However, since stop words and punctuations have been removed from this data, we used a punctuation restoration model² and then apply NLTK sentence segmentation for segmenting the job description.

B Hyperameters Tuning

Table 6 shows the hyperparameters chosen for grid search. The search was based on the validation performance (mAP@25 or MRR). The AdamW optimizer was used with a linear warm-up for 10% of the training steps with a batch size of 16. Every model was trained for ten epochs on the JTG-Synonym task and five epochs on the job normalization task. We calculated R@5, R@10, MRR, and mAP@25 using the trec_eval Python package (Van Gysel and de Rijke, 2018). Every experiment was conducted using PyTorch (Paszke et al., 2019) and NVIDIA RTX 3090GPU with 24GB memory.

All models ended up with a temperature parameter, τ , of 0.05. For our skill-based approach, an initial learning rate of $3e-5$ and $1e-5$ was used in the JTG-Synonym and job normalization task, respectively. For our JD-based approach, an initial learning rate of $1e-5$ and $3e-5$ were used in the JTG-Synonym and job normalization task, respectively. The number of transformer layers used in the job description aggregator was 4.

	List of values
Learning Rate	[$1e-4$, $1e-5$, $3e-5$, $5e-5$, $3e-6$, $5e-6$]
Temperature (τ)	[0.1, 0.05, 0.01]
# layers of JD aggregator	[1, 2, 4]

Table 6: Hyperparameters and their values used in grid search.

²<https://huggingface.co/felflare/bert-restore-punctuation>

C Detailed Dataset Description

C.1 JTG-Jobposting

Dataset statistics of the JTG-Jobposting dataset are shown in Table 3. Overall, the dataset statistics are very similar to the Mycarrersfuture.sg dataset, though the former has a much higher number of distinct skills. This is because skill in the JTG-Jobposting dataset is more relaxed compared to Mycarrersfuture.sg where a predefined set of skills is available. Some data examples from JTG-Jobposting are shown in Figure 4. It consists of the following fields:

- **"Position_Name"**: a job title of the job posting
- **"Skill_Hashtag"**: the skill tags assigned by the recruiter.
- **"Job_Description"**: a job description for the job posting.

C.2 JTG-Synonym

The JTG-Synonym is a synonym list that includes different variants of the same job title. We split the list into validation and test of size 2,000 and 4,420, respectively. Thai and English titles were kept separated. Examples are shown in Figure 5. The statistics of the JTG-Synonym are shown in Table 7.

	Total Query	Query			Candidate Pool	
		Thai	English	Code-Switching	Thai	English
Validation	2,000	1,044	928	28	7,762	8,033
Test	4,420	2,261	2,103	56	16,905	17,684

Table 7: The statistics of query and candidate pool in validation and test set for JTG-Synonym. "Code-Switching" refers to the job title that contains both Thai and English.

D Extended Implementation Detail

This subsection further describes our method and competing approaches. The model weight and inference code are available at <https://github.com/SLSCU/JD-aggr-network>.

D.1 Linear Probing

We applied linear probing by freezing the whole model except for the last feedforward layer. The performance of linear probing was evaluated on another set of JTG-Jobposting data to ensure no overlapping with the training data. The objective of the task is to predict the appropriate skills given

Job title	Recruiter Annotation	Keyword Matching	Skill Extraction	Job Description
Installation Manager (exp. M&E railway/เพศชาย) "BTS ท่าอากาศยานสุวรรณภูมิ"	MRT, BTS, รถไฟฟ้า, ระบบรถไฟฟ้า, railway system, M&E, Electrical, Mechanical, Engineering	'Control', 'coordinate', 'Design', 'drawing', 'drawings', 'installation', 'program', 'site', 'statement'	sql, 'solid work', 'python', 'java', 'electrical', 'ออกแบบ', 'design', 'เขียนแบบ', 'mechanical', 'autocad'	- To coordinate, supervise and control site installation activities. - To ensure installation tasks are delivered timely. - Able to produce detailed design drawings, installation method of statements and schedule with Primavera program.
นักวิเคราะห์ข้อมูล [Translated: Data analyst]	Excel, Python, SQL, Statistics	'การขาย', 'การขายและการตลาด', 'การ ตลาด', 'กิจกรรม', 'ขาย', 'ขายและการตลาด', 'ข้อมูล', 'ตลาด', 'ประสาน', 'ประสานงาน', 'ลูกค้า', 'วิเคราะห์', 'วิเคราะห์ ข้อมูล', 'สนับสนุน'	'iells', 'งานขาย', 'toefl', 'การ ตลาด', 'marketing', 'english language', 'ภาษาอังกฤษ', 'sale', 'english', 'toeic', 'sales'	- ประสานงาน รวบรวมข้อมูล วิเคราะห์ข้อมูลในการ ขายและการตลาด [Translated: Coordinate, gather data, and analyze sales and marketing data] - นำเสนอเพื่อสนับสนุนกิจกรรมให้ตรงตาม ความต้องการลูกค้า [Translated: Make a presentation to support operations corresponding to customer needs]

Figure 3: Examples of annotated skills that are not explicitly mentioned in the job description but are presented in the recruiter annotation (red highlights).

```
{
  "Position_Name": "นักพัฒนาโปรแกรม Java,  
[Translated: Java Developer]"
  "Skill_Hashtag": [ "java", "weblogic", "oracle", "ajax", "พัฒนาโปรแกรม", "เขียนโปรแกรม"],  
[Translated: "java", "weblogic", "ajax", "Programing Development", "Programing"]
  "Job_Description": "- พัฒนาเว็บไซต์และโปรแกรมด้วย Java บน Weblogic  
[Translated: Developing websites and being able to program Java on Weblogic.]  
- สามารถพัฒนาเว็บไซต์ด้วย HTML, CSS, JavaScript  
[Translated: Developing websites with HTML, CSS and JavaScript.]  
- มีประสบการณ์การเชื่อมต่อและใช้งาน Oracle ",  
[Translated: Has experiences in connecting and using Oracle.]
  "Company_Name": ABC ,
}
```

Figure 4: An example of job posting in the JTG-Jobposting dataset.

the job titles. The dataset contains 6,861 training samples and 2,000 testing samples consisting of 157 classes (skills). Due to the sparsity of skill labeling, we evaluated the performance using top-10 accuracy. The experiment was conducted using job title embeddings from three sources: our skill-based method with recruiter annotation, JD-based method, and XLM-R without fine-tuning.

D.2 Skill Extraction Model

The model used for skill extraction is mUSE (Yang et al., 2020) followed by a 2-layer MLP with cross-entropy loss as the objective. During inference, we select the top 10 scores as candidates for representing the skill tags for each job posting. The model was trained on another separate set of the JTG-Jobposting dataset containing 12,240 samples, totaling 35,107 skills, that do not overlap with the

original JTG-Jobposting dataset.

D.3 Comparison against other methods

Since the performance of the skill-based method of the JTG-Jobposting was not available, we reimplemented the baselines using the following configurations:

- **JobBERT**: We followed (Decorte et al., 2021) by randomly selected five samples from a distribution defined by the frequency distribution of skills in the training corpus, raised to the power of 3/4 for training using the skip-gram technique. We used a batch size of 64 and a learning rate of 5e-6.
- **Doc2VecSkill**: We reimplemented this baseline by aggregating the skill set in each job title, and then Doc2vec was used to convert

```

{
  "เจ้าหน้าที่ปฏิบัติการสินเชื่อ" :
  {
    "en": [ "Installment Loan Officer", "Loan Operation Officer",
    "th": [ "จนท.สินเชื่อ", "พนักงานสินเชื่อ"],
  },
  "Chief Accountant Specialist" :
  {
    "en": [ "Accountant Supervisor", "General Accounting Supervisor"],
    "th": [ "ซูเปอร์ไวเซอร์ผู้ควบคุมสินทรัพย์บัญชี", "ซูเปอร์ไวเซอร์ผู้ควบคุมงานบัญชี"],
  },
  "ฝ่ายจัดงาน Event":
  {
    "en": [ "Event Organizer", "Event Executive", "Event Planner", "Special Event Officer"],
    "th": [ "เจ้าหน้าที่กิจกรรมการตลาด", "เจ้าหน้าที่จัดกิจกรรม"],
  },
  },
  .....
}

```

Figure 5: An example of synonym in the JTG-Synonym evaluation dataset.

this set into auxiliary skill embeddings. The Doc2vec model was trained for 100 epochs with a dimension size of 768. Then, we matched the auxiliary skill embeddings with their embedded job titles using the cosine similarity loss. We used a batch of 64 and a learning rate of 3e-5.

E Qualitative Results

Figure 6 shows examples of attention maps in the first attention layer from the job description aggregator. It was found that the aggregator could correctly attend to sections with high importance and ignore the sentences unrelated to the job title. For example, in row 2 of Figure 6, the sentences "managing document of the developed software" and "Being studious, self-taught, responsible, and self-improvement" were mostly ignored while the sentence "Design and develop CRM Web Application ..." was strongly attended. The observation also held even when the first line was not the most informative sentence (example number 3 and 4).

F Design choice for our skill-based method

We explored different design choices for combining multiple skills to train our skill-based method. These included averaging, maximizing the output of embedded skills, and concatenating the skill list. The results in Table 8 suggest that concatenation, our final choice, performed best.

	JTG-Synonym	Job Normalization
	R@10 ↑	R@10 ↑
Skill-Based		
Concat	60.62 (± 0.32)	56.45 (± 0.18)
Max	52.63 (± 0.45)	53.19 (± 0.45)
Mean	58.51 (± 0.54)	55.83 (± 0.20)

Table 8: The comparison of different design choices of the skill-based method.

G Design choice for the sentence representation aggregation

We explored different approaches for representing the sentence embedding from a sequence of tokens by comparing token aggregation using an average with directly using the [CLS] token. It was found that averaging the tokens in the sentence yielded marginal performance improvement over the usage of the [CLS] token.

	Job Normalization
	R@10 ↑
Mean	57.09 (± 0.39)
[CLS]	56.84 (± 0.16)

Table 9: JD-based (Document Level): Ablation of different method: "Mean" denotes averaging each tokens (words), "[CLS]" denotes using the [CLS] token.

Job Title	Job Description
Admin Officer	<p>'1. ตรวจสอบสรุปออเดอร์ที่เข้ามาแต่ละเดือน', (52.58) [Translated : Check monthly order summaries.]</p> <p>'2. ทำอินพุต เอาพุต สรุปยอดการผลิต', (14.95) [Translated : Input-output for production summary]</p> <p>'3. เบิกของ', (15.94) [Translated : Pick up goods]</p> <p>'4. ทำรีพอร์ต', (2.43) [Translated : Prepare reports]</p> <p>'5. สามารถใช้excel ในระดับดีมาก', (1.06) [Translated : Proficient in Excel]</p> <p>'6. งานอื่นๆ ที่ได้รับมอบหมาย', (1.38) [Translated : Other assigned tasks]</p> <p>'7. สามารถย้ายที่อยู่ได้ (นายจะเปิดโรงงานใหม่ที่ไทรน้อย นนทบุรี)' (6.03) [Translated : Able to relocate (the company plans to open a new factory in Nonthaburi)]</p> <p>Special token -CLS (5.63)</p>
Web Developer (C#)	<p>'1) วางแผน ออกแบบ และพัฒนา ระบบ Web Application ในเชิง CRM (Customer Relationship Management) และ POS (Point Of Sale) ทั้งในส่วนของหน้างาน (Front- End)', (52.18) [Translated : Plan, design, and develop Web Application systems focusing on CRM (Customer Relationship Management) and POS (Point Of Sale), both in the front-end]</p> <p>'2) แก้ไขและปรับปรุงปัญหาต่างๆ ที่เกิดขึ้นกับตัวระบบ', (16.68) [Translated : Troubleshoot and resolve various system issues]</p> <p>'3) พัฒนาระบบตามแนวทางและข้อบังคับต่างๆ ที่ทางบริษัทกำหนด รวมไปถึงการนำเอาหลักการการพัฒนาซอฟต์แวร์ใหม่ๆ มาประยุกต์เพื่อให้ระบบทำงานได้อย่างถูกต้องและมีประสิทธิภาพที่สุด', (22.73) [Translated : Develop systems according to company guidelines and regulations, including incorporating new software development principles to ensure accurate and efficient system operation]</p> <p>'4) จัดการงานด้านเกี่ยวกับเอกสารต่างๆ ของระบบที่พัฒนาขึ้น', (6.01) [Translated : Manage documentation-related tasks for the developed systems]</p> <p>'5) ชยัน สามารถเรียนรู้ได้ด้วยตนเอง มีความรับผิดชอบ พัฒนาตัวเองอย่างสม่ำเสมอ' (2.36) [Translated : Diligent, self-learning, responsible, and committed to continuous self-improvement]</p> <p>Special token - CLS (0.01)</p>
Sales Manager (50K-60K)	<p>'- Salary 50,0 to 60,0 THB Industry Trading Firm Location BRT Wat Dorkmai Work Type Mon', (3.8)</p> <p>'- Fri 8:30?17:30 Job description', (0.65)</p> <p>'- Management sale control and monitor overall for sale department', (16.69)</p> <p>'- Present products and approach service to the customers', (4.79)</p> <p>'- Maintain relationship with clients by providing support, information and guidance, researching and recommending new opportunities, profit and service improvements', (15.62)</p> <p>'- Support basic technical of the products to the customers', (8.58)</p> <p>'- Sales for dealer, plant maker and engineering company', (15.20)</p> <p>'- Visit Client with Sales agency', (5.04)</p> <p>'- Take care existing client', (9.25)</p> <p>'- Sales for own products (Pump for construction site) in Thai', (17.12)</p> <p>'- Guidance and training of subordinates' (2.85)</p> <p>Special token - CLS (0.45)</p>
Quality Engineer	<p>'- Job description' (0.46)</p> <p>'- Support new program development process - product quality inspection.' (2.10)</p> <p>'- Conduct product inspection for new line installation and new model set up' (23.09)</p> <p>'- Create W/I and inspection check sheet.' (2.42)</p> <p>'- Perform Product Analysis (part claim -Warranty claim line functional claim analysis)' (18.48)</p> <p>'- Planning for development potential of subordinator and Drive team to reach the target.' (16.16)</p> <p>'- Support team and other parties for problem solving and smooth operation with all QA job concern.'(19.95)</p> <p>'- Manage manpower for support in product schedule plan.' (7.69)</p> <p>'- Manage about quality control in process (Incoming , Machine , Assy , Packing Out going)' (8.94)</p> <p>Special token - CLS (0.66)</p>

Figure 6: Attention scores from the CLS token of sentences extracted from the first attention layer of the job description aggregator. Darker red means higher attention scores and vice versa.