

Learning Job Titles Similarity from Noisy Skill Labels

Rabih Zbib, Lucas Alvarez Lacasa, Federico Retyk, Rus Poves, Juan Aizpuru, Hermenegildo Fabregat, Vaidotas Šimkus*, and Emilia García-Casademont*

Avature Machine Learning
machinelearning@avature.net

Abstract. Measuring semantic similarity between job titles is an essential functionality for automatic job recommendations. This task is usually approached using supervised learning techniques, which requires training data in the form of equivalent job title pairs. In this paper, we instead propose an unsupervised representation learning method for training a job title similarity model using noisy skill labels. We show that it is highly effective for tasks such as text ranking and job normalization.

Keywords: Semantic Text Similarity · Information Retrieval · Neural Networks · Job Title Normalization · Text Ranking

1 Introduction

With the significant growth of online platforms for job postings and applications, intelligent recommendation systems have become a necessity for both applicants and recruiters. Measuring semantic similarity between job titles is an essential functionality of these systems, whether to recommend suitable job openings to candidates, or vice versa. Job title similarity can be used as the sole measure for relevance, or more generally, as a component for computing an overall score between jobs and candidates, together with other information such as education, skills, and location. Recent work on this task goes beyond explicit lexical similarity and uses machine learning models that represent the semantics of the job title [13]. It takes the usual approach of learning semantic similarity by training a Siamese network using training data in the form of similar pairs [14]. But such labeled data is often hard to obtain in adequate amounts and quality. Skills, on the other hand, have recently become a major focus of the talent management industry [3], as they are highly informative of candidate abilities and job requirements.

In this paper, we present *Job Similarity Training*, a new unsupervised technique for training a job title encoder for similarity tasks in two stages. We first compute a dense-vector *auxiliary* embedding on sets of skills associated with a job title. Second, we train a neural job title encoder to mimic those auxiliary embeddings. The encoder maps a new job title into a dense vector that can be

* Work done while at Avature.

used for semantic similarity, text retrieval, and job normalization tasks. However, assigning accurate and comprehensive skills to candidate profiles or job descriptions is a specialized and time-consuming process, and automatic extraction of skills using a machine learning model also requires high-quality training data. We instead use *noisy* skills extracted from job listings and resumes with simple string matching.

We conduct a series of experiments that compare our method with alternative approaches, and show that it is highly effective for retrieval and normalization. The contributions of this paper are summarized as follows:

- *Job Similarity Training*, a new unsupervised training procedure for learning a job title encoder trained on noisy skills.
- A new evaluation dataset that we release as a benchmark for measuring the performance of job title encoders in a text ranking task.
- A set of experiments that show the effectiveness of the proposed unsupervised procedure for job title semantic similarity tasks, where we show that it outperforms a series of baselines, including the current state-of-the-art [1].
- Additional experiments, presented in the Appendices, which help to analyze the work from different perspectives: in Appendix A we study the efficiency of the training procedure, in Appendix B we study the performance of our method in multilingual and cross-lingual text retrieval tasks, and in Appendix C we study the extent to which the resulting job title encoder model captures information about the associated skills.

2 Related Work

Deep learning approaches for text ranking typically involve an encoder that converts the input text (query or document) into a dense, fixed-size representation, and then uses a similarity function to rank results. Huang et al. [4] use an MLP-based architecture for the encoder, and cosine similarity for ranking. They train the MLP using a supervision signal from a downstream retrieval task. Palangi et al. [15] train an LSTM-based encoder to predict similarity of sentences, and apply it to text ranking within a web search engine, with user clicks as supervision. A similar approach, but using BERT [2] with an attention layer, was explored by Humeau et al. [5]. For a detailed review of the use of BERT in information retrieval (IR) architectures, we refer the reader to Chapter 5 of the survey by Lin et al. [10]. Our method is based on a similar scheme, except that we use an unsupervised training procedure instead of supervised learning.

There has recently been an increase in work relevant to tasks in the talent management domain. Javed et al. [6] normalize job titles using a hierarchy of successively more specific classifiers. Neculoiu et al. [14] present a supervised, end-to-end deep learning approach to the same problem. They train a character-based BiLSTM encoder to map job titles to a fixed-length vector, such that the distance between the vectors is a measure of the job title semantic similarity. The encoder is trained as a Siamese network with supervised contrastive loss.

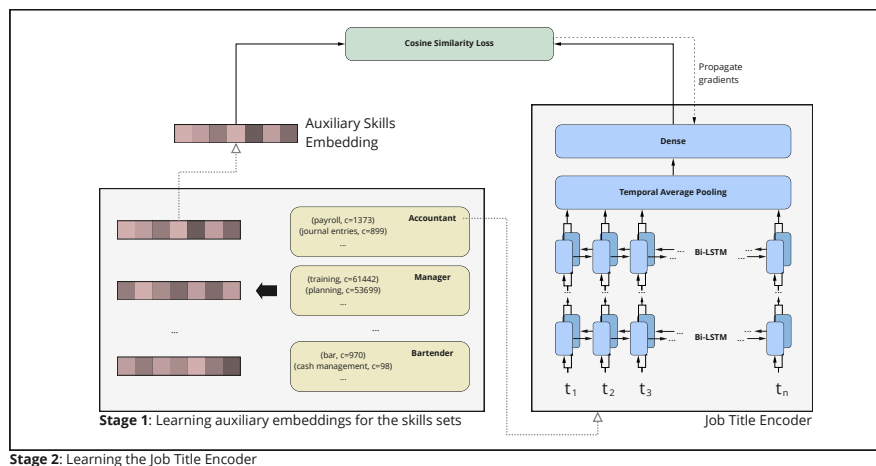


Fig. 1: Overview of *Job Similarity Training*, exemplified with a BiLSTM encoder. In **Stage 1**, we obtain an auxiliary representation for each job title using the set of related noisy skills. These embeddings only use the skills for each job (i.e. the job title text is not used so far). In **Stage 2**, we train the job title encoder to map the job titles to their skills auxiliary representations. The trained encoder is used for downstream tasks, such as IR or job normalization.

Zhang et al. [18] and Yamashita et al. [17] present methods for learning job title representations by using graphs to model the relationship between jobs in candidate career paths, combining different ways of encoding the job title into a multi-view graph-based representation.

Finally, JobBERT by Decorte et al. [1] is most similar to our approach. They use a pre-trained BERT language model to extract features for each token in the job title, and aggregate those features into a fixed-length vector using a soft attention layer. To train the job title encoder (including fine-tuning BERT) without a labeled dataset, they use the skills associated with each job title. Inspired by the Negative Sampling technique used for Skip-gram [12], they train the encoder as a classifier that predicts whether or not a particular skill corresponds to a job title. In our work, we use a different training objective and different neural architecture design. We compare the two approaches experimentally in Section 4.

3 Proposed Method

We next present the paper’s main contribution: **Job Similarity Training**, an unsupervised training procedure for learning semantic representations of job titles. We train a neural network encoder of job titles in two stages. We first learn an embeddings-based auxiliary representation for each job position using its noisy skills (and requiring no human-labeled supervision), and then train the encoder to map the job title text to its auxiliary representation (see details in Figure 1).

3.1 Data Collection and Preprocessing

The input to the entire training procedure is a dataset \mathcal{D}_{raw} of pairs (j, \mathbf{s}_j) , where j is the job title text, \mathbf{s}_j is a set of skills $s \in \mathcal{S}$, and \mathcal{S} is a predefined skills vocabulary. In our experiments, the dataset is constructed from English job postings and from the work experience section of anonymized resumes. The skills are extracted from the description text using simple string matching, producing a noisy version of the skill set associated with each job.

Next, we generate the dataset $\mathcal{D}_{\text{merged}}$ by aggregating the noisy skill sets that correspond to the same job title. This aggregation results in a multiset $\mathbf{s}_j^+ = \{s_1: m_1, \dots, s_{n_j}: m_{n_j}\}$ where $s_i \in \mathcal{S}$ is a skill and m_i is its count. Aggregating skill counts across common job titles reduces the noise since a skill that is extracted multiple times is more likely to be correct.

3.2 Skills Auxiliary Representation

In training Stage 1, we use the dataset $\mathcal{D}_{\text{merged}}$ to produce an auxiliary representation for each job title. We use the Doc2vec model [9] to directly compute a fixed-length vector \mathbf{e}_j for each skill multiset \mathbf{s}_j^+ , considering it as a document and treating individual skills as words. The output is a dataset \mathcal{D}_{aux} of pairs (j, \mathbf{e}_j) , with j a job title and \mathbf{e}_j its auxiliary representation. These representations are then used as synthetic targets in the second training stage.

In the experiments, we use the Distributed Bag of Words (PV-DBOW) version of Doc2vec. PV-DBOW obtained better results than other approaches, including the Distributed Memory (PV-DM) version of Doc2vec as well as the average vector of individual skill embeddings learned with Word2vec [11]. Full details of the comparison are excluded for space considerations.

3.3 Job Title Encoder

In Stage 2 we train an encoder η that maps the job title text to its final representation, motivated by the assumption that the auxiliary representations learned in the first stage capture important semantic information about the job titles j . Using the dataset \mathcal{D}_{aux} , the encoder is trained to minimize the cosine distance between $\eta(j)$ and the corresponding auxiliary representation \mathbf{e}_j .

It is important to observe that, while the encoder η is trained using information about the noisy skill sets (via the auxiliary representations \mathbf{e}_j), the only information visible to the encoder during inference is the job title text. The encoder is thus forced to learn to infer information about the related skills. When provided with a new, unseen job title j , the model is able to generalize and infer about job-to-skills relationship. Similar to JobBERT [1], we also recognize that skills are useful information for building representations of job titles. But our unsupervised training procedure is markedly different: while they use a contrastive learning objective to train the encoder, we train it to mimic the representations \mathbf{e}_j obtained from the skills. We compare the two approaches in Section 4.

Method		MAP	P@5	P@20
Text-based Retrieval				
Model	Training Method			
Okapi BM25	Trained on $\mathcal{D}_{\text{merged}}$	0.2754	0.5067	0.3062
BERT	None (no fine-tuning)	0.1556	0.3124	0.1871
BiLSTM	Negative Sampling	0.6428	0.7581	0.5376
BERT	Negative Sampling	0.6011	0.7238	0.5152
BiLSTM	Job Similarity Training	0.6814	0.7790	0.5781
BERT	Job Similarity Training	0.7077	0.7829	0.5929
Skill-based Retrieval				
TF-IDF (Noisy Test Skills)		0.3319	0.5481	0.3135
Doc2vec (Noisy Test Skills)		0.1031	0.1675	0.1204
TF-IDF (Gold Standard Test Skills)		0.7880	0.8376	0.6668
Doc2vec (Gold Standard Test Skills)		0.7126	0.7446	0.5921

Table 1: Performance of different unsupervised models on the text ranking evaluation set. Mean Average Precision (MAP), Precision at 5, and Precision at 20 are shown for each condition. Bold font style corresponds to the method achieving the best results.

4 Experiments

We next study the proposed approach through a series of experiments in two downstream tasks: text ranking and job normalization. As described in Section 3, for the training data we extract job titles and skills from job postings and anonymized resumes. The resulting raw dataset \mathcal{D}_{raw} contains 44 million samples, and the merged set $\mathcal{D}_{\text{merged}}$ 8.5 million sample pairs. The set of skills we use contains 5,600 skills covering a variety of industries.

4.1 Text Ranking

Text ranking involves ranking a set of text documents in a corpus given a query and a measure of relevance. In the context of a job title recommendation system, both the documents and the queries are short text. The corpus job titles are ranked by first computing a vector representation using the encoder $\eta(j)$, and then computing the cosine distance between the query and the job title vectors.

We use a new job title similarity dataset¹ for evaluation, consisting of 2,724 job titles from different industries (105 of which are used as queries). Each query/corpus-item pair is labeled for binary relevance after adjudicating two independent human annotations². The main experimental results are in Table 1. Following common practice, we evaluate models using the Mean Average Precision (MAP) of the output of the ranked lists, as well as Precision at 5 and 20. We use the `trec_eval` software library to compute these metrics³.

Text-based Retrieval. The *Job Similarity Training* experiment use two different architectures for the encoder η : a BiLSTM encoder trained from scratch, with

¹ https://github.com/rabihzbib/jobtitlesimilarity_dataset

² With an inter-annotator agreement measured at 86%

³ https://github.com/usnistgov/trec_eval

SentencePiece [8] for tokenization and average pooling to aggregate the token representations; and a pre-trained BERT-base encoder [2] using its built-in pooling mechanism. Both architectures have a linear Dense layer on top to produce a vector with the same dimensions as the auxiliary skills representations.

We compare *Job Similarity Training* to two simple baselines. The first is the well-known ranking function Okapi BM25 [7, 16], trained and evaluated using the job title text and ignoring the skills. The second baseline is a pre-trained BERT-based language model with no fine-tuning. Table 1 shows that both baselines obtain MAP scores well below the models trained with either Job Similarity Training or Negative Sampling. These results support the hypothesis that using skill-related information is important to train the encoder effectively.

Job Similarity Training is also compared to a similar unsupervised alternative: the Negative Sampling scheme proposed by Decorte et al. [1] to train JobBERT. In this case, the encoder is trained to predict whether a skill is related to a job title. Results show that our method outperforms Negative Sampling for both architectures. For a fair comparison, the same two encoder architectures are used for both methods.

Skill-based Retrieval. A natural question in the context of this work is how effective would it be to rank jobs by looking only at their set of skills, and not the job title text. To answer this, we use the skills of the test set jobs to build two vector representations: discrete, high-dimensional TF-IDF skill vectors, and Doc2vec embeddings of the skill sets (i.e., the synthetic target \mathbf{e}_j used to train $\eta(j)$). In both cases, the skill vectors are used directly to compute the cosine distance⁴. We use two versions of the test job skills:

- *Test Job Noisy Skills.* For each job title in the test set, we extract a skill set from its description text, consistent with how noisy skills are extracted for training.
- *Gold Standard Skills.* For each job title in the test set, Talent Management experts define a skill set. This version is intended to measure the effect of the quality of the skills on retrieval results.

We report on two experiments for each of these two sets, using Doc2vec and TF-IDF vectors. The performance of the Noisy Skill-based baselines is significantly lower than that of our method, showing that using the noisy skills directly to represent the job at inference time is not adequate. The Gold Standard baselines outperform our methods, but they use the gold-standard skill set to represent each test jobs. From a practical perspective, the difference between the Noisy and the Gold Standard baselines shows that high-quality skills can be directly used for retrieval, but they are often not available in practice. With noisy skills, however, the performance is significantly degraded. Our method leverages the noisy skills to learn a semantic representation of the job titles. The model *generalizes* to unobserved job titles, as shown by the significantly better

⁴ In this paper’s method, as in JobBERT [1], skills are not needed at inference time. The test job skills are used here for comparison only.

Model	Training Method	MRR	P@5	P@10
BERT	Decorte et al. [1]	0.3092	0.3865	0.4604
BiLSTM	Job Similarity Training	0.3007	0.3955	0.4760
BERT	Job Similarity Training	0.3414	0.4595	0.5400

Table 2: Performance of different unsupervised models on Decorte [1]’s test set. Mean Reciprocal Rank (MRR) and Precision at 5 and 10 of the micro averages are shown.

performance that the job title encoder achieves compared to the Noisy Skills baselines.

4.2 Job Normalization

The second downstream task is job normalization, which consists of mapping an input job title to one element in a set of normalized job titles. Here, we also encode the raw input job as well as the normalized job titles, and select the one corresponding to the smallest cosine distance. We use the test set of Decorte et al. [1], which consists of 15,463 raw job titles (as typically found on online job boards) and their corresponding normalized job title. 2,675 ESCO occupations⁵ are used as the normalized set. Table 2 shows that the Job Similarity Training models also outperform the results reported by JobBERT for job normalization.

5 Conclusion

In this paper, we presented *Job Similarity Training*, a new unsupervised technique for training job title encoders to learn semantic information. We evaluated it on two downstream tasks: text ranking, using a new evaluation dataset that we release as part of this work, and job normalization, using an existing benchmark.

Our experiments show that the new procedure outperforms various baselines in both tasks, including JobBERT (the previous state-of-the-art). Furthermore, they show that, while the quality of the skill-to-job mapping is a critical factor in the performance of systems that directly use skills for understanding the semantics of a job title, our method can leverage noisy skills to obtain good results.

We also presented additional experiments (included in the Appendices) where we explore the efficiency of the method when compared to JobBERT, its performance in multilingual and cross-lingual scenarios, and the encoder’s ability to capture information about skills related to a job title.

⁵ https://esco.ec.europa.eu/en/classification/occupation_main

References

1. Decorte, J.J., Hautte, J.V., Demeester, T., Develder, C.: JobBERT: Understanding Job Titles through Skills. In: FEAST, ECML-PKDD 2021 Workshop (2021)
2. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 4171–4186. Association for Computational Linguistics (Jun 2019)
3. Djumalieva, J., Sleeman, C.: An Open and Data-driven Taxonomy of Skills Extracted from Online Job Adverts. Tech. Rep. ESCoE DP-2018-13, Economic Statistics Centre of Excellence (ESCoE) (Aug 2018)
4. Huang, P.S., He, X., Gao, J., Deng, L., Acero, A., Heck, L.: Learning deep structured semantic models for web search using clickthrough data. In: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management. p. 2333–2338. CIKM '13, Association for Computing Machinery, New York, NY, USA (2013)
5. Humeau, S., Shuster, K., Lachaux, M.A., Weston, J.: Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In: International Conference on Learning Representations (2020)
6. Javed, F., Luo, Q., McNair, M., Jacob, F., Zhao, M., Kang, T.S.: Carotene: A job title classification system for the online recruitment domain. In: 2015 IEEE First International Conference on Big Data Computing Service and Applications. pp. 286–293 (2015)
7. Jones, K.S., Walker, S., Robertson, S.E.: A probabilistic model of information retrieval: development and comparative experiments. *Information Processing and Management* **36**, 779–808 (November 2000)
8. Kudo, T., Richardson, J.: SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. pp. 66–71. Association for Computational Linguistics, Brussels, Belgium (Nov 2018)
9. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: Xing, E.P., Jebara, T. (eds.) Proceedings of the 31st International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 32, pp. 1188–1196. PMLR, Beijing, China (22–24 Jun 2014)
10. Lin, J., Nogueira, R., Yates, A.: Pretrained Transformers for Text Ranking: BERT and Beyond. Morgan & Claypool Publishers (2021)
11. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space (2013)
12. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems* **26** (2013)
13. Mitra, B., Craswell, N.: Neural models for information retrieval. arXiv preprint arXiv:1705.01509 (2017)
14. Neculoiu, P., Versteegh, M., Rotaru, M.: Learning text similarity with Siamese recurrent networks. In: Proceedings of the 1st Workshop on Representation Learning for NLP. pp. 148–157. Association for Computational Linguistics, Berlin, Germany (Aug 2016)

15. Palangi, H., Deng, L., Shen, Y., Gao, J., He, X., Chen, J., Song, X., Ward, R.: Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* **24**(4), 694–707 (2016)
16. Robertson, S., Zaragoza, H.: The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval* **3**(4), 333–389 (Apr 2009)
17. Yamashita, M., Li, Y., Tran, T., Zhang, Y., Lee, D.: Looking further into the future: Career pathway prediction. In: *In Proceedings of the First International Workshop on Computational Jobs Marketplace* (2022)
18. Zhang, D., Liu, J., Zhu, H., Liu, Y., Wang, L., Wang, P., Xiong, H.: Job2vec: Job title benchmarking with collective multi-view representation learning. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. pp. 2763–2771 (2019)

A Analysis of Training Procedure Efficiency

In this section, we analyze the difference between Job Similarity Training and Negative Sampling in terms of efficiency. In order to do this, we fine-tune a BERT-small encoder with each unsupervised method (Similarity Training, Negative Sampling), and periodically measure the performance of these models in the text ranking task. The encoder architecture, the training hyperparameters, and the number of steps per epoch are identical in both cases. The encoder trained with Job Similarity Training uses $\mathcal{D}_{\text{merged}}$ as its training dataset. The one trained with Negative Sampling uses positive samples from the dataset \mathcal{D}_{raw} and generate negative samples according to that method’s definition. Figure 2 compares the performance of both methods under these conditions. We aggregate the learning curves of three replications for each method.

Besides obtaining a superior performance at convergence, Job Similarity Training achieves a better MAP score much earlier in the training process. In other words, given a fixed amount of forward passes, our method performs significantly better in the downstream task. This is particularly important when fine-tuning large Transformer-based language models like BERT.

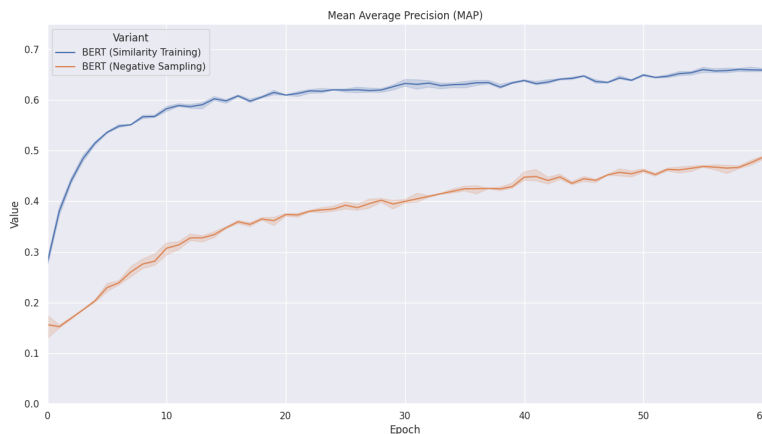


Fig. 2: Performance, as measured in a downstream text ranking task, for different model snapshots during the unsupervised training procedures: Job Similarity Training and Negative Sampling.

B Multi- and Cross-lingual Text Retrieval

Although in our experiments, skills are extracted from English job descriptions, those skills represent language-independent concepts. This implies that the auxiliary representations that we use to train the job title encoder η are also language-independent. So all that is needed to build an encoder in a new language is to

translate the job titles j included in each sample from the dataset \mathcal{D}_{aux} to obtain a new dataset $\mathcal{D}_{\text{aux}}^{\mathbf{L}}$ for the target language \mathbf{L} . The corresponding model $\eta_{\mathbf{L}}$ can thus be created without manual data effort and without repeating the initial data collection and preprocessing workload, which might be harder in the case of low-resource languages. Also, the language independence of the semantic space can be further exploited to build multilingual models by training the job title encoder η on the combined data from multiple languages. The resulting model can be used for cross-lingual retrieval.

In this experiment, we take the original English dataset \mathcal{D}_{aux} and use Machine Translation (MT) to translate around 1 million job titles j to German and 2 million job titles j to French, in order to obtain $\mathcal{D}_{\text{aux}}^{\text{DE}}$ and $\mathcal{D}_{\text{aux}}^{\text{FR}}$ respectively. These are used to train either monolingual or bilingual models η with the same BiLSTM architecture explored in Section 4.1, and trained with the proposed Similarity Training procedure.

To evaluate these models, we translate the evaluation set using human translators. The rationale is that, using MT for the training set, one can transfer the model to another language quickly, even if original data in that language is not available. On the other hand, producing high-quality human translations of the evaluation set is a better reflection of the real scenario, where users input native job titles rather than machine-translated ones.

The results in Table 3 show that this approach produces a viable model at no extra human effort in translation or labeling. The results for the German and the French monolingual models (MAP 0.5545 and 0.5355, respectively) are much higher than the Noisy-skills baselines from Table 1. The encoder-based English models have better performance, as they were trained on 8.5 million samples, while the German models were trained on 1 million and the French models on 2 million.

We also experiment with bilingual models, combining data from two languages to train a French-English encoder and a German-English encoder. We observe that the bilingual model can be used for cross-lingual retrieval —when either the query or the corpus titles are in English— without loss in performance, providing evidence to the assumption that the information that can be captured from the skill set is correlated with the underlying, language-independent job position concept, and therefore is transferable across languages. The cross-lingual retrieval capability is achieved with no penalty in monolingual retrieval quality, since we see a small difference of less than 1% between monolingual and bilingual models for retrieval when the query and the corpus are in the same language (first and second lines in each section of the table). In practice, adding English to models in another language is a useful capability, since job titles in general are likely to include English terms.

Training	Q_{Test}	C_{Test}	MAP	P@5	P@20
DE	DE	DE	0.5545	0.7000	0.4808
DE+EN	DE	DE	0.5476	0.6885	0.4745
DE+EN	DE	EN	0.5942	0.6712	0.5202
DE+EN	EN	DE	0.5545	0.7010	0.4890
FR	FR	FR	0.5355	0.6692	0.4750
FR+EN	FR	FR	0.5437	0.6923	0.4707
FR+EN	FR	EN	0.5908	0.6827	0.4933
FR+EN	EN	FR	0.5577	0.6857	0.4867

Table 3: Monolingual and bilingual model results on either monolingual or cross-lingual text retrieval tasks. For all experiments, we use the *Doc2vec (Noisy Skills) + Encoder* variant of the training procedure. Language codes: German (DE), English (EN), and French (FR).

C Predicting Skills from the Encoded Job Title

Lastly, we discuss the results of a different kind of experiment, where the encoded job titles are used to predict related skills. The aim is to further show that the encoded job title representation captures information about the skills. This functionality is also useful to suggest possible skills for a user when creating a job description.

A consequence of the proposed training procedure is that the representation spaces of the encoded job titles and the auxiliary embeddings are compatible, in the sense that measuring the distance between an encoded job $\eta(j)$ and the skill auxiliary embeddings is meaningful. The same argument applies for the individual skill embeddings produced by Skip-gram and word2vec.

In this experiment, we use the job title encoder η to encode the job titles from the evaluation set, and for each job title we retrieve the N skills whose individual skills embeddings are closest in the shared representation space. Finally, we measure the precision of the predicted top- N skills using the gold standard set of skills for that job title (the gold standard used for the skill-based baselines in Section 4.1). We obtain the following results: P@5 is 0.4577, P@10 is 0.3854, and P@20 is 0.3127. Considering that generating a set of elements from a large vocabulary—in this case, 5,600 skills—is a difficult problem, even for humans, these results show that the encoder η captures information about the skills effectively.

Furthermore, to allow for a qualitative assessment of the extent to which the model generalizes to new job titles, we show the predicted skills for two job titles that are not included in the training set.

Director of Communications:

Public Relations, Social Media, Media Relations,
Campaigns, Strategy, Writing, Editing

Construction and Building Inspector:

Construction, Inspection, Materials,
Contractors, Structures, Plumbing,
Concrete, Regulations, Completion, Welding

The second example is particularly interesting since the position involves the intersection of two concepts (*Construction/Building* and *Inspector*), and the predicted skills reflect that.