

KOMPETENCER: Fine-grained Skill Classification in Danish Job Postings via Distant Supervision and Transfer Learning

Mike Zhang^{*◇} Kristian Nørgaard Jensen^{*◇} Barbara Plank^{◇♣}

[◇]Department of Computer Science, IT University of Copenhagen, Denmark

[♣]Center for Information and Language Processing (CIS), LMU Munich, Germany

{mikz, krnj}@itu.dk bplank@cis.uni-muenchen.de

Abstract

Skill Classification (SC) is the task of classifying job competences from job postings. This work is the first in SC applied to Danish job vacancy data. We release the first Danish job posting dataset: KOMPETENCER (*en*: competences), annotated for nested spans of competences. To improve upon coarse-grained annotations, we make use of The European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al. (2014)) taxonomy API to obtain fine-grained labels via distant supervision. We study two setups: The zero-shot and few-shot classification setting. We fine-tune English-based models and RemBERT (Chung et al., 2020) and compare them to in-language Danish models. Our results show RemBERT significantly outperforms all other models in both the zero-shot and the few-shot setting.

Keywords: Skill Classification, Distant Supervision, Transfer Learning, Domain Adaptive Pretraining, Job Postings

1. Introduction

Job Posting data (JPs) is emerging on a variety of platforms in big quantities, and can provide insights on labor market skill set demands and aid job matching (Balog et al., 2012). *Skill Classification* (SC) is to classify competences (i.e., hard and soft skills) necessary for any occupation from unstructured text or JPs.

Several works focus on Skill Identification (Jia et al., 2018; Sayfullina et al., 2018; Tamburri et al., 2020). This is to classify whether a skill occurs in a sentence or job description. However, continuing the pipeline, there is little work in further categorizing the identified skills by leveraging taxonomies such as ESCO. Another limitation is the scope of language, where all previous work focus on English job postings. This hinders in particular local job seekers from finding an occupation suitable to their specific skills within their community via online job platforms.

In this work, we look into the Danish labor market. We introduce KOMPETENCER, a novel Danish job posting dataset annotated on the *span-level* for nested *Skill* and *Knowledge* Components (SKCs) in job postings. We do not directly annotate for the fine-grained taxonomy codes from e.g., ESCO, but rather annotate more generic spans of SKCs (Figure 2), and then exploit the ESCO API to bootstrap fine-grained SKCs via distant supervision (Mintz et al., 2009) and create “silver” data for skill classification. Our proposed distant supervision pipeline is denoted in Figure 1.

Recently, Natural Language Processing has seen a surge of several transfer learning methods and architecture which help improve state-of-the-art significantly on several tasks (Peters et al., 2018; Howard and Ruder, 2018; Radford et al., 2018; Devlin et al., 2019). In this work, we explore the benefits of zero-shot cross-lingual transfer learning with English BERT_{base} (De-

Distant Supervision for Skill Classification




Figure 1: **Pipeline for Fine-grained Danish Skill Classification.** We propose a distant supervision pipeline, where we have identified spans of skills and knowledge. We query the ESCO API and fine-tune a model on the distantly supervised labels.

vlin et al., 2019) and a BERT_{base} that we continuously pretrain (Han and Eisenstein, 2019; Gururangan et al., 2020) on 3.2M English JP sentences and test on Danish and compare it to in-language models: Danish BERT and our domain-adapted Danish BERT model on 24.5M Danish JP sentences. We analyze the zero-shot transfer of English to Danish SC. Last, we experiment with few-shot training: We fine-tune a multilingual model (Chung et al., 2020) on English JPs with a few Danish JPs and show how zero-shot transfer compares to training on a small amount of in-language data.

Contributions ① We release KOMPETENCER,¹ the first Danish Skill Classification dataset with distantly supervised fine-grained labels using the ESCO taxonomy. ② We furthermore present experiments and analysis with in-language Danish models vs. a zero-shot cross-lingual transfer from English to Danish with domain-adapted BERT models. ③ We target a few-shot learning setting with a multilingual model trained on both English and a few Danish JPs.

^{*}The authors contributed equally to this work.

¹<https://github.com/jjzha/kompetencer>

↓ Statistics, Language →	ENGLISH (EN)	DANISH (DA)
# Posts	391	60
# Sentences	14,538	1,479
# Tokens	232,220	20,369
# Skill Spans	6,576	665
# Knowledge Spans	6,053	255
\bar{x} Skill Span	3.97	3.71
\bar{x} Knowledge Span	1.80	1.73
\tilde{x} Skill Span	4	3
\tilde{x} Knowledge Span	1	1
Skill [90%]	[1, 9]	[1, 9]
Knowledge [90%]	[1, 5]	[1, 4]
Silver fine-grained labels	✓	✗
Gold fine-grained labels	✗	✓

Table 1: **Statistics of Annotated Dataset.** We report the total number of JPs across languages and their respective number of sentences, tokens, and SKCs. Below, we show the mean length of SKCs (\bar{x}), median length of SKCs (\tilde{x}), and the 90th percentile of length [90%] starting from length 1. We also indicate the type of labels in both sets (silver or gold labels). The EN set is larger than the DA split.

2. KOMPETENCER Dataset

2.1. Skill & Knowledge Definition

There is an abundance of competences and there have been large efforts to categorize them. The European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al. (2014)) taxonomy is the standard terminology linking skills, competences and qualifications to occupations. The ESCO taxonomy mentions three categories of competences: *Knowledge*, *skill*, and *attitudes*. ESCO defines knowledge as follows:

“Knowledge means the outcome of the assimilation of information through learning. Knowledge is the body of facts, principles, theories and practices that is related to a field of work or study.”²

For example, a person can acquire the Python programming language through learning. This is denoted as a *knowledge* component and can be considered generally a *hard skill*. However, one also needs to be able to apply the knowledge component to a certain task. This is known as a *skill* component. ESCO formulates it as:

“Skill means the ability to apply knowledge and use know-how to complete tasks and solve problems.”³

In ESCO, the *soft skills* are referred to as *attitudes*. ESCO considers attitudes as skill components:

“The ability to use knowledge, skills and personal, social and/or methodological abilities, in work or study situations and professional and personal development.”⁴

²ec.europa.eu/esco/portal/escopedia/Knowledge

³ec.europa.eu/esco/portal/escopedia/Skill

⁴data.europa.eu/esco/skill/A

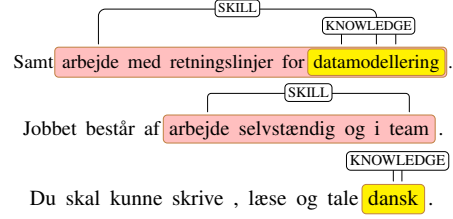


Figure 2: **Examples of Skills and Knowledge Components.** Annotated samples of passages in varying Danish job postings. SKCs can be nested as shown in the first example.

To sum up, hard skills are usually referred to as *knowledge* components, and applying these hard skills to something is considered a *skill*. Then, soft skills are referred to as *attitudes*, these are part of skill components. There has been no work, to the best of our knowledge, in annotating skill and knowledge components in JPs.

2.2. Dataset Statistics

Both the English and Danish data comes from a large job platform with various types of JPs.⁵ The English JPs are from Zhang et al. (2022). In Table 1, we show the statistics of both the annotated English and Danish data split. We note that the number of English JPs is larger than the Danish split. For Danish, there are fewer knowledge spans proportional to English. Apart from this, both the English and Danish JPs follow a similar trend in terms of statistics. The mean length of skills and knowledge (\bar{x}) is slightly shorter for Danish, 3.97 vs. 3.71 and 1.80 vs. 1.73 respectively. The median length of skills (\tilde{x}) is one token shorter for Danish. However, we note again that the length of skills can vary substantially, ranging from 1–9 for both languages. Then, for knowledge components this ranges from 1–5 and 1–4 for English and Danish respectively. The similarity in statistics shows the consistency of annotations, which we elaborate on in the next section.

Figure 2 shows some examples of the annotated SKCs. “Samt arbejde med retningslinjer for datamodellering” (en: “As well as working with guidelines for data modeling”), shows a nesting example: “datamodellering” shows a knowledge component (i.e., something that one can learn), and the skill is to apply it. “Jobbet består af arbejde selvstændig og i team” (en: The job consists of working independently and in a team) indicates an *attitude* as “working independently or in a team” is a social ability. We furthermore consider languages a knowledge component, as one can acquire the language through schooling. Overall, the classification of the spans could be a short sentence (i.e., ≤ 9 tokens) or single token classification.

⁵We release the annotated spans in <https://github.com/jjzha/kompetencer/tree/master/data>




Figure 3: **Label Distribution of Distantly Supervised Labels.** In the top and middle barplot we show the fine-grained label distribution of the English training and development split respectively. The splits follow a similar distribution. For the Danish training split on the bottom, there is a large increase of A1 labels, which indicate more *attitude*-like skills. All splits have a larger fraction of the label S1, which encapsulates communicative skills. Explanations of labels are given in Table 3 (Section 11, Appendix).

2.3. Annotations

Skill Identification Annotations We annotate with the annotation guidelines denoted in Zhang et al. (2022) used on the English data split to identify the SKCs in a JP. There are around 57.5K tokens (approximately 4.6K sentences, in 101 job posts) that was used to calculate agreement on. The annotations were compared using Cohen’s κ (Fleiss and Cohen, 1973) between pairs of annotators, and Fleiss’ κ (Fleiss, 1971), which generalizes Cohen’s κ to more than two concurrent annotations. We consider two levels of κ calculations: **TOKEN** is calculated on the token level, comparing the agreement of annotators on each token (including non-entities) in the annotated dataset. **SPAN** refers to the agreement between annotators on the exact span match over the surface string, regardless of the type of named entity, i.e., we only check the position of tag without regarding the type of the named entity. The observed agreement scores over the three annotators is between 0.70–0.75 Fleiss’ κ for both levels of calculation which is considered a *substantial agreement* (Landis and Koch, 1977). Then, for the Danish data split, we use the same guidelines as for English. Here, we consider one annotator that annotates for the SKCs.

Fine-grained Annotations Currently, our proposed dataset consists of identified SKCs. To obtain fine-grained labels of each span, we explore distant supervision using the ESCO API, where the setup is broadly depicted in Figure 1. The annotated spans are queried to the API, then via Algorithm 1, we determine whether the obtained SKC is “relevant” or not via Levenshtein distance matching (Levenshtein, 1966). In addition, we determine the quality of the distant supervised labels by human evaluation. We manually check each of the annotated spans to its obtained label from the ESCO API.

Algorithm 1 Getting the best match for a skill in the ESCO API using Levenshtein distance

```

procedure FETCHSKILL(Skill, Type) ▷ Find Skill
in the ESCO API
   $X \leftarrow$  Top-100 query results from ESCO
   $X \leftarrow \{ \text{typeof}(x) = \textit{Type} : x \in X \}$ 
   $d \leftarrow \infty$ 
   $r \leftarrow \text{None}$ 
  for  $x \in X$  do
     $D \leftarrow \text{levenshtein}(x, \textit{Skill})$ 
    if  $D = 0$  then
      return  $x$  ▷ Perfect match
    else if  $D < d$  then
       $r \leftarrow x$ 
       $d \leftarrow D$ 
    end if
  end for
  return  $r$  ▷ Best match based on Levenshtein
distance
end procedure

```

After checking a subset 2,622 English labels — without correcting — and its distantly supervised labels, we obtain 41.3% accuracy on the correctness of the distantly supervised labels. We note that across all 9,473 labels in the original English training and development data (details of train/dev/test splits in Section 3.2), a total of 7.4% is unidentified by the ESCO database, and is thus labeled by K99 from ESCO in the resulting train and development data here. For the Danish data, we obtain 70.4% accuracy on the training set and 20.2% is missing, albeit the Danish training set only contains 138 SKCs. For the Danish test set, we correct the distantly supervised labels to create a gold test set.




Figure 4: **Experimental Pipeline for Fine-grained Danish Skill Classification.** Read from left to right, we start with each respective dataset for English (EN) and Danish (DA). We obtain the labels from the ESCO API and train for each language split two models: For EN, these are (1) BERT_{base} and (2) JobBERT. For DA, these are (3) DaBERT and (4) DaJobBERT. The Danish data is split into 10/50 train/test and the English data in to 290/51/50 train/dev/test JPs. The Danish models are fine-tuned on the Danish train set, and use *no* in-language development set (i.e., English dev.). In the end all models are applied to the Danish and English test set separately.

Here, 14.1% was initially correct and 23.5% missed a label. In Figure 3, we show the distantly supervised fine-grained label distribution of the English training and development set split, and the Danish training split. The following labels: 0000, K?, and S? are artifacts of querying the ESCO API (i.e., unidentified skills). We did not employ any post-processing and left them as is. We presumed they would not influence the model significantly as their numbers are low.

3. Methodology

In the current setting, we have annotated spans of SKCs. We extract the spans from the JPs and query the ESCO API to obtain silver labels. We formulate this task as a text classification problem. We consider a set of JPs \mathcal{D} , where $d \in \mathcal{D}$ is a set of extracted spans (*and not full sentences*) with the i^{th} span $\mathcal{X}_d^i = \{x_1, x_2, \dots, x_T\}$ and a target class $c \in \mathcal{C}$, where $\mathcal{C} = \{S^*, K^*\}$. The labels S^* and K^* depend on the distantly supervised ESCO taxonomy code (e.g., S4: Management Skills,⁶ K2: Arts and Humanities⁷). The goal of this task is then to use \mathcal{D} to train an algorithm $h: \mathcal{X} \mapsto \mathcal{C}$ to accurately predict skill tags by assigning an output label c for input \mathcal{X}_d^i .

3.1. Encoders

As baseline for Danish SC, we consider a Danish BERT (DaBERT) encoder.⁸ Following Gururangan et al. (2020), we continuously pretrain DaBERT on 24.5M Danish JP sentences for *one* epoch, we name this **Da-JobBERT**.⁹ To test zero-shot performance from En-

glish to Danish for SC, we use BERT_{base} (Devlin et al., 2019) and a domain-adapted BERT_{base} model on 3.2M JP sentences, namely JobBERT (Zhang et al., 2022). We assume that domain-adapted models like JobBERT and DaJobBERT would improve SC as the “domain” is the same.

Multilingual Encoder We explore whether using a multilingual encoder would benefit the classification of skills for Danish in a low-resource setting. For the experiments we use RemBERT (Chung et al., 2020), it has recently shown to outperform mBERT (Devlin et al., 2019) on several tasks. All models are using a final Softmax layer for the classification of spans.

3.2. Experimental Setup

Our detailed experimental setup is shown in Figure 4. We start with 391 English and 60 Danish job postings (Table 1) annotated with spans of SKCs. The spans are then queried to the ESCO API (Figure 1). We split the English data into 290 train (9,472 SKCs), 51 dev (1,577 SKCs), and 50 JPs for test (1,578 SKCs), and for the Danish data we split this into 10 JPs (138 SKCs) for training and 50 JPs for test (782 SKCs). For the label distribution we refer back to Figure 3 (excl. test).

We fine-tune BERT_{base} and JobBERT on the spans in 290 English JPs. Next, we fine-tune DaBERT and DaJobBERT on the 10 Danish JPs. For RemBERT, we fine-tune in three ways: Only on English, only on Danish, and on both English and Danish together. For all setups, we choose the model with the best score on the English dev. set. As pointed out by Artetxe et al. (2020): Pure unsupervised cross-lingual transfer should not use any cross-lingual signal by definition. As our attention is on Danish, we do not use any Danish labeled training data *nor* dev. data in the zero-shot setting. All models in the end will be tested on

⁶<http://data.europa.eu/esco/skill/S4>

⁷<http://data.europa.eu/esco/isced-f/02>

⁸<https://huggingface.co/Maltheb/danish-bert-botxo>

⁹<https://huggingface.co/jjzha/dajobbert-base-cased>




Figure 5: **Performance of Models on English and Danish.** We test seven setups on several splits of data: English development (**DEV (EN)**), English test set (**TEST (EN)**), and Danish test set (**TEST (DA)**). Reported is the weighted macro-F1. The whiskers indicate each respective standard deviation of runs on five random seeds. Left side of the black vertical line indicates a full zero-shot setting on **TEST (DA)**, on the right shows the few-shot setting on the same test set. With respect to the models, language abbreviation in brackets (e.g., **BERT (EN)**) indicates what it has been fine-tuned on. Exact numbers including significance testing are noted in Table 4 (Section 12, Appendix).

the held-out 50 English and Danish JPs separately.¹⁰ In summary, we have three setups: (1) Fine-tuned on English JPs only (BERT, JobBERT, RemBERT), (2) fine-tuned on Danish JPs only (DaBERT, DaJobBERT, RemBERT), and (3) fine-tuned on both English and Danish JPs (RemBERT). We consider (1) a zero-shot setting, while (2) and (3) do have access to some Danish training data, hence this is a few-shot setting. Throughout the experiments, we use the MACHAMP (v0.3) toolkit (van der Goot et al., 2021) for classification. All reported results are the average over five runs with different random seeds on weighted macro-F1.

4. Analysis of Results

We show the experimental results in Figure 5. Plotted is the weighted macro-F1 of all three setups with seven models and their corresponding standard deviation on the English development set, English test set, and the Danish test set. All models left of the black vertical line are the zero-shot setup, applied to Danish. On the right, these models are in the few-shot setting, this is due to the model having access to some target language training data (DA).

Performance Zero-shot Setting For the models trained on English only (BERT, JobBERT, and RemBERT (EN)) when applied to the English development set, all three models perform similarly. They achieve around 0.63–0.64 weighted macro-F1 with little standard deviation: $BERT_{base}$ 0.628 ± 0.004 , JobBERT 0.628 ± 0.006 , and RemBERT (EN) 0.629 ± 0.003 weighted macro-F1. Similarly for the English test set: $BERT_{base}$ 0.632 ± 0.007 , JobBERT 0.644 ± 0.006 ,

and RemBERT (EN) 0.637 ± 0.007 weighted macro-F1, where JobBERT significantly outperforms all other models (details in Section 12).

It is a tacit that the English-based models perform better than the baseline (DaBERT) on English, both dev. and test. Conversely, the English-based models perform poorly on the Danish test set: $BERT_{base}$ 0.038 ± 0.008 and JobBERT 0.063 ± 0.005 weighted macro-F1. However, given a multilingual encoder (RemBERT) only trained on English, gives a significant gain in zero-shot performance (0.354 ± 0.021) with little standard deviation and significantly outperforms the other zero-shot setting models including the target-language baseline (DaBERT). We strongly suspect this is due to Danish being included in the pretraining data of RemBERT.

Performance Few-shot Setting Apart from RemBERT (EN+DA) having access to English data, all other models fine-tuned on Danish perform poorly on English dev. and test. The performance of RemBERT (DA) is slightly better than the best performing Danish-only model DaJobBERT (0.098 ± 0.040 vs. 0.096 ± 0.024 weighted macro-F1 on English test), where our intuition again goes to the pretraining data.

For DA test, DaBERT is a strong baseline, achieving 0.199 ± 0.058 weighted macro-F1 with little Danish training data. RemBERT (DA) did not result in significant gains having pretrained on multiple languages and another intuition could be that this is a result of negative transfer (Rosenstein et al., 2005). Then, DaJobBERT performs better than DaBERT on the Danish test set: 0.395 ± 0.021 weighted macro-F1. Note that we conducted domain adaptive pretraining from the DaBERT checkpoint on 24.5M Danish JP sentences for one epoch with the Masked Language Modeling objec-

¹⁰The English test set contains silver labels (distantly supervised), while the Danish test set is human corrected (gold).




Figure 6: **Confusion Matrix of RemBERT (EN)**. We show the confusion matrix of the zero-shot setting with RemBERT (EN). On the diagonal are the correctly predicted labels. Most of the “confusion” is with respect to the labels that encompass the larger fraction of the test set: A1: Attitudes and S1: Communication, collaboration and creativity.

tive. This shows that in-language *and* in-domain pre-training is beneficial for this specific task of SC.

Combining Training Data Last, giving RemBERT all training data (English and Danish) results in substantial improvement over all other models in the zero-shot and few-shot setting alike: 0.472 ± 0.014 , which significantly outperforms all other models on Danish test. Henceforth, it is helpful to have a bit of target-language training data for higher resulting performance.

Is Domain Adaptive Pretraining Worth It? In light of the results, domain adaptive pretraining shows its benefit for both English and Danish fine-tuning. Specifically for Danish, from the baseline (DaBERT), the improvement is close to 0.2 weighted macro-F1 with DaJobBERT. The domain adaptive pretraining took ~ 35 hours, using 4 GPUs, to pass once over the unlabeled data (24.5M Danish JP sentences). The largest gain is obtained with combining both English and Danish training data: The improvement is around 0.27 weighted macro-F1. However, the 391 EN and 60 DA JPs took around two months of non-stop annotating. In short, there is a trade-off between continuous pretraining on unlabeled text and annotating: (1) Domain adaptive pretraining gives short-term gains with little costs, but there needs to be enough unlabeled data in the right domain. (2) Annotating extra data results in larger gains long-term, but there is more costs involved.

Analysis of Predictions In Figure 6, we show the confusion matrix of the best performing zero-shot model on the test set of the best run and investigate what the model does not predict correctly. In the matrix, the model mostly confuses the label A1, which relates to *attitudes* and gets predicted as S1: Communication, collaboration and creativity. There could be some overlap between these labels as for example the skill “effektiv” (*en*: efficient/effective). This is officially labeled as an attitude by ESCO, but a grey area is that “effective” could relate to “creativity”.

There is also a small cluster of confusion from S1–4. These are rather distinct classes of skills. For example, S4 means management skills. A specific example is “fagligt velfunderet” (*en*: professionally sound), this could be an attitude. This is all hard to determine since there is no context around the skill. Overall, there is some confusion between the skills when taken out of context. We leave the exploration of fine-grained skill classification *with context* for future work.

Qualitative Analysis Distant Supervision We analyze the label selection method and the missing labels in the English dataset as mentioned in Section 2.3. We find that the missing labels in the English data is predominantly coming from technical skills. We found that the missing spans are mostly knowledge components in the form of technologies used today by developers, such as ReactJS, Django, AWS etc. This lack of coverage could either be due to specificity or the ever-growing list of technologies. In ESCO, there are several technologies that are listed (e.g., NoSQL, Drupal, WordPress to name a few), but there are also a lot missing (e.g., TensorFlow, Data Science, etc.).

5. Related Work

Many works focus on the identification of skills in job descriptions, i.e., whether a sentence contains a skill or not (Sayfullina et al., 2018; Tamburri et al., 2020) or what the necessary skills are inferred from an entire job posting (Bhola et al., 2020). We instead identified the SKCs manually in the job descriptions on the sentence-level, as this gives us the highest quality of identified SKCs. Furthermore, there are several works in fine-grained SC (i.e., categorize the skills), but mostly focus on English job descriptions. A straightforward approach is to do exact matching with a predefined list of skills (Malherbe and Aufaure, 2016; Papoutsoglou et al., 2017; Sibarani et al., 2017) or do a frequency analysis of skills, cluster them by hand and attach a more general category to them e.g., Gardiner et al. (2018). Some works have used the ESCO taxonomy directly (Boselli et al., 2018; Giabelli et al., 2020). For example, Boselli et al. (2018) classified both titles and description for its most suitable ISCO (Elias, 1997) code (what ESCO is partially based on). However, they only gave one label to each data point (i.e., full job posting), which is unrealistic as most occupations require multiple competences.

Overall, to the best of our knowledge, there seems to be little to no work in directly classifying the identified SKC to a specific ESCO label. In addition, this work is the first of its kind doing this for Danish JPs.

6. Conclusion

We present a novel skill classification dataset for competences in Danish: KOMPETENCER.¹¹ In addition, we transform the coarse-grained human annotated spans to more fine-grained labels via distant supervision with the ESCO API. Our human evaluation shows that the distantly supervised labels give a signal of correctly annotated spans, where we achieve 41.3% accuracy on a large English label subset, and 70.4% accuracy on the Danish dev set, and 14.1% accuracy on the Danish test set. We manually correct the Danish test set with the correct labels from ESCO to create a gold annotated set and keep the English labels as is, and thus silver labels.

Furthermore, domain adaptive pretraining helps to improve performance on the task specifically for English. The best performance is achieved with RemBERT on both the zero-shot setting (0.354 ± 0.021 weighted macro-F1) and few-shot setting (0.472 ± 0.014 weighted macro-F1), where they significantly outperform the other models. The strong performance is likely due to the pretraining data that contains both Danish and English.

Last, since the annotations are on the token-level, this work can be extended to, for example, sequence labeling. We hope this dataset initiates further research in the area of skill classification.

7. Acknowledgements

We thank the NLPnorth group for feedback on an earlier version of this paper—in particular Elisa Bassignana and Max Müller-Eberstein for insightful discussions. We would also like to thank the anonymous reviewers for their comments to improve this paper. Last, we also thank NVIDIA and the ITU High-performance Computing cluster for computing resources. This research is supported by the Independent Research Fund Denmark (DRF) grant 9131-00019B.

8. Bibliographical References

Artetxe, M., Ruder, S., Yogatama, D., Labaka, G., and Agirre, E. (2020). A call for more rigor in unsupervised cross-lingual learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7375–7388.

¹¹We release the Danish anonymized raw data and annotations of the parts with permissible licenses from a governmental agency which is our collaborator. Links to our English data can be found at <https://github.com/kris927b/SkillSpan>. For anonymization, we perform it via manual annotation of job-related sensitive and personal data regarding Organization, Location, Contact, and Name following the work by Jensen et al. (2021).

- Balog, K., Fang, Y., De Rijke, M., Serdyukov, P., and Si, L. (2012). Expertise retrieval. *Foundations and Trends in Information Retrieval*, 6(2–3):127–256.
- Bender, E. M. and Friedman, B. (2018). Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Bhola, A., Halder, K., Prasad, A., and Kan, M.-Y. (2020). Retrieving skills from job descriptions: A language model based extreme multi-label classification framework. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5832–5842, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Bonferroni, C. (1936). Teoria statistica delle classi e calcolo delle probabilità. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62.
- Boselli, R., Cesarini, M., Mercorio, F., and Mezzananza, M. (2018). Classifying online job advertisements through machine learning. *Future Generation Computer Systems*, 86:319–328.
- Chung, H. W., Févry, T., Tsai, H., Johnson, M., and Ruder, S. (2020). Rethinking embedding coupling in pre-trained language models. *arXiv preprint arXiv:2010.12821*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dror, R., Shlomov, S., and Reichart, R. (2019). Deep dominance - how to properly compare deep neural models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2785, Florence, Italy. Association for Computational Linguistics.
- Elias, P. (1997). Occupational classification (isco-88): Concepts, methods, reliability, validity and cross-national comparability. Technical report, OECD Publishing.
- Fleiss, J. L. and Cohen, J. (1973). The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Gardiner, A., Aasheim, C., Rutner, P., and Williams, S. (2018). Skill requirements in big data: A content analysis of job advertisements. *Journal of Computer Information Systems*, 58(4):374–384.
- Giabelli, A., Malandri, L., Mercorio, F., and Mezzan-

- zanica, M. (2020). Graphlmi: A data driven system for exploring labor market information through graph databases. *Multimedia Tools and Applications*, pages 1–30.
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., and Smith, N. A. (2020). Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Han, X. and Eisenstein, J. (2019). Unsupervised domain adaptation of contextualized embeddings for sequence labeling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4238–4248, Hong Kong, China. Association for Computational Linguistics.
- Howard, J. and Ruder, S. (2018). Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339.
- Jensen, K. N., Zhang, M., and Plank, B. (2021). De-identification of privacy-related entities in job postings. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 210–221, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Jia, S., Liu, X., Zhao, P., Liu, C., Sun, L., and Peng, T. (2018). Representation of job-skill in artificial intelligence with knowledge graph analysis. In *2018 IEEE Symposium on Product Compliance Engineering-Asia (ISPCE-CN)*, pages 1–6. IEEE.
- Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- le Vrang, M., Papantoniou, A., Pauwels, E., Fannes, P., Vandestein, D., and De Smedt, J. (2014). Esco: Boosting job matching in europe with semantic interoperability. *Computer*, 47(10):57–64.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.
- Malherbe, E. and Aufaure, M.-A. (2016). Bridge the terminology gap between recruiters and candidates: A multilingual skills base built from social media and linked data. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 583–590. IEEE.
- Mintz, M., Bills, S., Snow, R., and Jurafsky, D. (2009). Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011.
- Papoutsoglou, M., Mittas, N., and Angelis, L. (2017). Mining people analytics from stackoverflow job advertisements. In *2017 43rd Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pages 108–115. IEEE.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Reimers, N. and Gurevych, I. (2018). Why comparing single performance scores does not allow to draw conclusions about machine learning approaches. *ArXiv preprint*, abs/1803.09578.
- Rosenstein, M. T., Marx, Z., Kaelbling, L. P., and Dietterich, T. G. (2005). To transfer or not to transfer. In *NIPS 2005 workshop on transfer learning*, volume 898, pages 1–4.
- Sayfullina, L., Malmi, E., and Kannala, J. (2018). Learning representations for soft skill matching. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 141–152.
- Sibarani, E. M., Scerri, S., Morales, C., Auer, S., and Collarana, D. (2017). Ontology-guided job market demand analysis: a cross-sectional study for the data science field. In *Proceedings of the 13th International Conference on Semantic Systems*, pages 25–32.
- Tamburri, D. A., Van Den Heuvel, W.-J., and Garriga, M. (2020). Dataops for societal intelligence: a data pipeline for labor market skills extraction and matching. In *2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)*, pages 391–394. IEEE.
- Ulmer, D. (2021). deep-significance: Easy and better significance testing for deep neural networks. <https://github.com/Kaleidophon/deep-significance>.
- van der Goot, R., Üstün, A., Ramponi, A., Sharaf, I., and Plank, B. (2021). Massive choice, ample tasks (MaChAmp): A toolkit for multi-task learning in NLP. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.
- Zhang, M., Jensen, K. N., Sonniks, S. D., and Plank, B. (2022). Skillspan: Hard and soft skill extraction from english job postings.

Appendix

9. Data Statement KOMPETENCER

Following Bender and Friedman (2018), the following outlines the data statement for KOMPETENCER:

PARAMETER	VALUE	RANGE
Optimizer	AdamW	
β_1, β_2	0.9, 0.99	
Dropout	0.2	0.1, 0.2, 0.3
Epochs	20	
Batch Size	32	
Learning Rate (LR)	1e-4	1e-3, 1e-4, 1e-5
LR scheduler	Slanted triangular	
Weight decay	0.01	
Decay factor	0.38	0.35, 0.38, 0.5
Cut fraction	0.2	0.1, 0.2, 0.3

Table 2: **Hyperparameters of MACHAMP.**

- A. CURATION RATIONALE: Collection of job postings in the English and Danish language for skill classification, to study the impact of skill changes from job postings.
- B. LANGUAGE VARIETY: The non-canonical data was collected from the StackOverflow job posting platform, an in-house job posting collection from our national labor agency collaboration partner (*which will be elaborated upon acceptance*), and web extracted job postings from a large job posting platform. US (en-US), British (en-GB) English, and Danish (da-DK) are involved.
- C. SPEAKER DEMOGRAPHIC: Gender, age, race-ethnicity, socioeconomic status are unknown.
- D. ANNOTATOR DEMOGRAPHIC: Three hired project participants (age range: 25–30), gender: one female and two males, white European and Asian (non-Hispanic). Native language: Danish, Dutch. Socioeconomic status: higher-education students. Female annotator is a professional annotator with a background in Linguistics and the two males with a background in Computer Science.
- E. SPEECH SITUATION: Standard American, British English or Danish is used in job postings. Time frame of the data is between 2012–2021.
- F. TEXT CHARACTERISTICS: Sentences are from job postings posted on official job vacancy platforms.
- G. RECORDING QUALITY: N/A.
- H. OTHER: N/A.
- I. PROVENANCE APPENDIX: The Danish job posting data is from our collaborators: The Danish Agency for Labour Market and Recruitment (STAR).

10. Reproducibility

We use the default hyperparameters in MACHAMP (van der Goot et al., 2021) as shown in Table 2. For more details we refer to their paper. For the five random seeds we use 3477689, 4213916, 6828303, 8749520, and 9364029. All experiments

with MACHAMP were ran on an NVIDIA® NVIDIA A100-SXM4 40GB GPU and an AMD® EPYC 7662 64-Core Processor.

11. Label Meaning

LABEL	SUBJECT	DEFINITION
0000	ARTIFACT	ARTIFACT
A1	Attitudes	Individual work styles that can affect how well someone performs a job.
A2	Values	Principles or standards of behavior, revealing one's judgment of what is important in life.
K00	Generic programmes and qualifications	Generic programmes and qualifications are those providing fundamental and personal skills education which cover a broad range of subjects and do not emphasise or specialise in a particular broad or narrow field.
K01	Education	NO-DEFINITION
K02	Arts and humanities	NO-DEFINITION
K03	Social sciences, journalism and information	NO-DEFINITION
K04	Business, administration and law	NO-DEFINITION
K05	Natural sciences, mathematics and statistics	NO-DEFINITION
K06	Information and communication technologies (icts)	NO-DEFINITION
K07	Engineering, manufacturing and construction not elsewhere classified	NO-DEFINITION
K08	Agriculture, forestry, fisheries and veterinary	NO-DEFINITION
K09	Health and welfare	NO-DEFINITION
K10	Services	NO-DEFINITION
K99	Field unknown	NO-DEFINITION
L1	Languages	Ability to communicate through reading, writing, speaking and listening in the mother tongue and/or in a foreign language.
S1	Communication, collaboration and creativity	Communicating, collaborating, liaising, and negotiating with other people, developing solutions to problems, creating plans or specifications for the design of objects and systems, composing text or music, performing to entertain an audience, and imparting knowledge to others.
S2	Information skills	Collecting, storing, monitoring, and using information; Conducting studies, investigations and tests; maintaining records; managing, evaluating, processing, analysing and monitoring information and projecting outcomes.
S3	Assisting and caring	Providing assistance, nurturing, care, service and support to people, and ensuring compliance to rules, standards, guidelines or laws.
S4	Management skills	Managing people, activities, resources, and organisation; developing objectives and strategies, organising work activities, allocating and controlling resources and leading, motivating, recruiting and supervising people and teams.
S5	Working with computers	Using computers and other digital tools to develop, install and maintain ICT software and infrastructure and to browse, search, filter, organise, store, retrieve, and analyse data, to collaborate and communicate with others, to create and edit new content.
S6	Handling and moving	Sorting, arranging, moving, transforming, fabricating and cleaning goods and materials by hand or using handheld tools and equipment. Tending plants, crops and animals.
S7	Constructing	Building, repairing, installing and finishing interior and exterior structures.
S8	Working with machinery and specialised equipment	Controlling, operating and monitoring vehicles, stationary and mobile machinery and precision instrumentation and equipment.
K?	ARTIFACT	ARTIFACT
S?	ARTIFACT	ARTIFACT

Table 3: **Definition of ESCO Labels.** Indicated are the definitions of the ESCO labels used in this work taken from the ESCO taxonomy. Artifacts of the ESCO API are K? and S?, and 0000, this means that no component was found.

MODEL	EN DEV	EN TEST	DA TEST
BERT _{base} (EN)	0.628±0.004	0.632±0.007	0.038±0.008
JobBERT (EN)	0.628±0.006	0.644±0.006*	0.063±0.005
RemBERT (EN)	0.629±0.003	0.637±0.007	0.354±0.021
DaBERT (DA)	0.088±0.013	0.076±0.012	0.199±0.058
DaJobBERT (DA)	0.101±0.024	0.096±0.024	0.395±0.021
RemBERT (DA)	0.116±0.052	0.098±0.040	0.166±0.141
RemBERT (EN+DA)	0.629±0.006*	0.643±0.006	0.472±0.014*

Table 4: **Exact Results on Splits.** Indicated are the exact results of the bar plots in Figure 5. Significance tested with Almost Stochastic Order (Dror et al., 2019) test with Bonferroni correction (Bonferroni, 1936). Bold indicates highest average weighted macro-F1 and asterisk indicates significance.

12. Exact Results from Plots

In Table 4, we show the exact results of the plots from Figure 5 on English dev, English test, and Danish test respectively. In addition, we do significance testing. Recently, the Almost Stochastic Order (ASO) test (Dror et al., 2019)¹² has been proposed to test statistical significance for deep neural networks over multiple runs. Generally, the ASO test determines whether a stochastic order (Reimers and Gurevych, 2018) exists between two models or algorithms based on their respective sets of evaluation scores. Given the single model scores over multiple random seeds of two algorithms \mathcal{A} and \mathcal{B} , the method computes a test-specific value (ϵ_{\min}) that indicates how far algorithm \mathcal{A} is from being significantly better than algorithm \mathcal{B} . When distance $\epsilon_{\min} = 0.0$, one can claim that \mathcal{A} stochastically dominant over \mathcal{B} with a predefined significance level. When $\epsilon_{\min} < 0.5$ one can say $\mathcal{A} \succeq \mathcal{B}$. On the contrary, when we have $\epsilon_{\min} = 1.0$, this means $\mathcal{B} \succeq \mathcal{A}$. For $\epsilon_{\min} = 0.5$, no order can be determined. We compared all pairs of models based on five random seeds each using ASO with a confidence level of $\alpha = 0.05$ (before adjusting for all pair-wise comparisons using the Bonferroni correction (Bonferroni, 1936)). Almost stochastic dominance ($\epsilon_{\min} < 0.5$) is indicated in Figure 7 over all the splits.

13. Confusion Matrix Few-Shot

In Figure 8, we show the confusion matrix of the best performing few-shot model on the test set of the best run and investigate what the model does not predict correctly. Dissimilar from Figure 6, we only seem some confusion in the small cluster of S1-4. Giving the model a few Danish JPs substantially improved the prediction of A1, which relates to *attitudes* and gets predicted as S1: Communication, collaboration and creativity.

¹²Implementation of Dror et al. (2019) can be found at <https://github.com/Kaleidophon/deep-significance> (Ulmer, 2021)

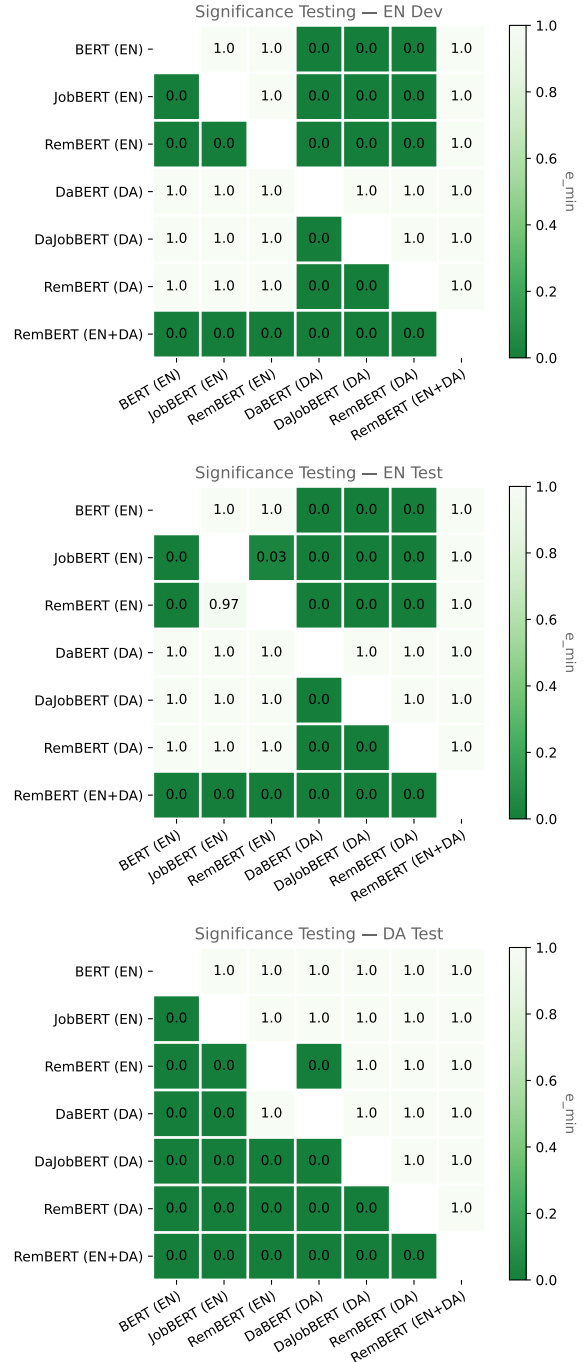


Figure 7: **Results Almost Stochastic Order.** ASO scores expressed in ϵ_{\min} . The significance level $\alpha = 0.05$ is adjusted accordingly by using the Bonferroni correction (Bonferroni, 1936). Almost stochastic dominance ($\epsilon_{\min} < 0.5$) is indicated in the colored boxes: On **EN TEST**, JobBERT is almost stochastically dominant over RemBERT (EN), with $\epsilon_{\min} = 0.03$.




Figure 8: **Confusion Matrix of RemBERT (EN+DA).** We show the confusion matrix of the few-shot setting with RemBERT (EN+DA). On the diagonal are the correctly predicted labels. There is less confusion in this model as compared to RemBERT (EN). We suspect the additional Danish data benefits the prediction of A1.