

SKILLSPAN: Hard and Soft Skill Extraction from English Job Postings

Mike Zhang^{*◇} Kristian Nørgaard Jensen^{*◇} Sif Dam Sonniks[◇] Barbara Plank^{◇♣}

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

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

{mikz, krnj, sifs}@itu.dk

bplank@cis.uni-muenchen.de

Abstract

Skill Extraction (SE) is an important and widely-studied task useful to gain insights into labor market dynamics. However, there is a lacuna of datasets and annotation guidelines; available datasets are few and contain crowd-sourced labels on the span-level or labels from a predefined skill inventory. To address this gap, we introduce SKILLSPAN, a novel SE dataset consisting of 14.5K sentences and over 12.5K annotated spans. We release its respective guidelines created over three different sources annotated for hard *and* soft skills by domain experts. We introduce a BERT baseline (Devlin et al., 2019). To improve upon this baseline, we experiment with language models that are optimized for long spans (Joshi et al., 2020; Beltagy et al., 2020), continuous pre-training on the job posting domain (Han and Eisenstein, 2019; Gururangan et al., 2020), and multi-task learning (Caruana, 1997). Our results show that the domain-adapted models significantly outperform their non-adapted counterparts, and single-task outperforms multi-task learning.

1 Introduction

Job markets are under constant development—often due to developments in technology, migration, and digitization—so are the skill sets required. Consequentially, job vacancy data is emerging on a variety of platforms in big quantities and can provide insights on labor market skill demands or aid job matching (Balog et al., 2012). SE is to extract the competences necessary from unstructured text.

Previous work in SE shows promising progress, but is halted by a lack of available datasets and annotation guidelines. Two out of 14 studies release their dataset, which limit themselves to crowd-sourced labels (Sayfullina et al., 2018) or annotations from a predefined list of skills on the document-level (Bhola et al., 2020). Additionally,

^{*}Equal contribution.

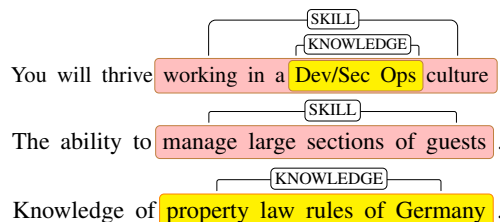


Figure 1: **Examples of Skills & Knowledge Components.** Annotated samples of passages in varying job postings. More details are given in Section 4.

none of the 14 previously mentioned studies release their annotation guidelines, which obscures the meaning of a competence. Job markets change, as do the skills in, e.g., the European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al., 2014) taxonomy (Section 3). Hence, it is important to cover for possible emerging skills.

We propose SKILLSPAN, a novel SE dataset annotated at the span-level for *skill* and *knowledge* components (SKCs) in job postings (JPs). As illustrated in Figure 1, SKCs can be nested inside skills. SKILLSPAN allows for extracting possibly undiscovered competences and to diminish the lack of coverage of predefined skill inventories.

Our analysis (Figure 2) shows that SKCs contain on average longer sequences than typical Named Entity Recognition (NER) tasks. Albeit we additionally study models optimized for long spans (Joshi et al., 2020; Beltagy et al., 2020), some underperform. Overall, we find specialized domain BERT models (Alsentzer et al., 2019; Lee et al., 2020; Gururangan et al., 2020; Nguyen et al., 2020) perform better than their non-adapted counterparts. We explore the benefits of domain-adaptive pre-training on the JP domain (Han and Eisenstein, 2019; Gururangan et al., 2020). Last, given the examples from Figure 1, we formulate the task as both as a sequence labeling and a multi-task learning (MTL) problem, i.e., training on both skill and knowledge components jointly (Caruana, 1997).





	Annotations	Approach	Size	Skill Type	(Baseline) Model(s)		
Kivimäki et al. (2013)	Document-level	Automatic	N/A	Hard	LogEnt., TF-IDF, LSA, LDA	✗	✗
Zhao et al. (2015)	Sentence-level	Automatic	N/A	Hard	Word2Vec	✗	✗
Javed et al. (2017)	Span-level	Skill Inventory	N/A	Both	Word2Vec	✗	✗
Jia et al. (2018)	Span-level	Automatic	21,158 JPs*	Hard	LSTM	✗	✗
Sayfullina et al. (2018)	Span-level	Crowdsourcing	4,863 Sent.	Soft	CNN, LSTM	✓	✗
Smith et al. (2019)	Span-level	Manual	100 JPs	Hard	Pattern Matching	✗	✗
Gugnani and Misra (2020)	Span-level	Domain Experts	~200 JPs	Hard	Word2Vec, Doc2Vec	✗	✗
Li et al. (2020)	Document-level	Proprietary	N/A	Hard	FastText	✗	✗
Shi et al. (2020)	Span-level	Proprietary	N/A	Hard	FastText, USE, BERT	✗	✗
Tamburri et al. (2020)	Sentence-level	Domain Experts	~3,000 Sent.	Both	BERT	✗	✗
Chernova (2020)	Span-level	Manual	100 JPs	Both	FinBERT	✗	✗
Bhola et al. (2020)	Document-level	Skill Inventory	20,298 JPs*	Hard	BERT	✓	✗
Smith et al. (2021)	Span-level	Manual	100 JPs	Hard	Pattern Match., Word2Vec	✗	✗
Liu et al. (2021)	Document-level	Crowdsourcing	N/A	Hard	GNN	✗	✗
This work	Span-level	Domain Experts	391 JPs	Both	(Domain-adapted) BERT	✓	✓

Table 1: **Contributions of Related Work.** We list the recent works of Skill Extraction. Note that (*) indicates labels that are automatically inferred from some source (e.g., a predefined skill inventory) and not *manually* annotated. With respect to the annotation approach, “Manual” indicates uncertainty whether they used domain experts or not. Also note that many works do not release their dataset with annotations () nor guidelines (). The list is inspired by Khaouja et al. (2021).

Contributions In this paper: ① We release SKILLSPAN, a novel skill extraction dataset, with annotation guidelines, and our open-source code.¹ ② We present strong baselines for the task including a new SpanBERT (Joshi et al., 2020) trained from scratch, and domain-adapted variants (Gururangan et al., 2020), which we will release on the HuggingFace platform (Wolf et al., 2020). To the best of our knowledge, we are the first to investigate the extraction of skills and knowledge from job postings with state-of-the-art language models. ③ We give an analysis on single-task versus multi-task learning in the context of skill extraction, and show that for this particular task single-task learning outperforms multi-task learning.

2 Related Work

There is a pool of prior work relating to SE. We summarize it in Table 1, depicting state-of-the-art approaches, level of annotations, what kind of competences are annotated, the modeling approaches, the size of the dataset (if available), type of skills annotated for, baseline models, and whether they release their annotations and guidelines.

As can be seen in Table 1, many works do *not release their data* (apart from Sayfullina et al., 2018 and Bhola et al., 2020) and **none release their annotation guidelines**. In addition, none of the previous studies approach SE as a span-level extraction task with state-of-the-art language models, nor

did they release a dataset of this magnitude with manually annotated (long) spans of competences by domain experts.

Although Sayfullina et al. (2018) annotated on the span-level (thus being useful for SE) and release their data, they instead explored several approaches to *Skill Classification*. To create the data, they extracted all text snippets containing *one* soft skill from a predetermined list. Crowdworkers then annotated the highlighted skill whether it was a soft skill referring to the candidate or not. They show that an LSTM (Hochreiter et al., 1997) performs best on classifying the skill in the sentence. In our work, we annotated a dataset three times their size (Table 2) for both hard *and* soft skills. In addition, we also extract the specific skills from the sentence.

Tamburri et al. (2020) classifies sentences that contain skills in the JP. The authors manually labeled their dataset with domain experts. They annotated whether a sentence contains a skill or not. Once the sentence is identified as containing a skill, the skill cited within is extracted. In contrast, we directly annotate for the span within the sentence.

Bhola et al. (2020) cast the task of skill extraction as a multi-label skill classification at the document-level. There is a predefined set of unique skills given the job descriptions and they predict multiple skills that are connected to a given job description using BERT (Devlin et al., 2019). In addition, they experiment with several additional layers for better prediction performance. We instead explore domain-adaptive pre-training for SE.

¹<https://github.com/kris927b/SkillSpan>

The work closest to ours is by Chernova (2020), who approach the task similarly with span-level annotations (including longer spans) but approach this for the Finnish language. It is unclear whether they annotated by domain experts. Also, neither the data nor the annotation guidelines are released. For a comprehensive overview with respect to SE, we refer to Khaouja et al. (2021).

3 Skill & Knowledge Definition

There is an abundance of competences and there have been large efforts to categorize them. For example, the The International Standard Classification of Occupations (ISCO; Elias, 1997) is one of the main international classifications of occupations and skills. It belongs to the international family of economic and social classifications. Another example, the European Skills, Competences, Qualifications and Occupations (ESCO; le Vrang et al., 2014) taxonomy is the European standard terminology linking skills and competences and qualifications to occupations and derived from ISCO. The ESCO taxonomy mentions three categories of competences: *Skill*, *knowledge*, 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 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

²<https://ec.europa.eu/esco/portal/escopedia/Knowledge>

³<https://ec.europa.eu/esco/portal/escopedia/Skill>

	↓ Statistics, Src. →	BIG	HOUSE	TECH	Total
Train	# Posts	60	60	80	200
	# Sentences	1,036	1,674	3,156	5,866
	# Tokens	29,064	36,995	56,549	122,608
	# Skill Spans	1,086	984	1,237	3,307
	# Knowledge Spans	439	781	2,188	3,408
	# Overlapping Spans	45	29	135	209
Development	# Posts	30	30	30	90
	# Sentences	783	1,022	2,187	3,992
	# Tokens	11,762	19,173	21,149	52,084
	# Skill Spans	469	525	545	1,539
	# Knowledge Spans	126	287	806	1,219
	# Overlapping Spans	12	17	32	61
Test	# Posts	36	33	32	101
	# Sentences	1,112	1,216	2,352	4,680
	# Tokens	14,720	21,923	20,885	57,528
	# Skill Spans	634	637	459	1,730
	# Knowledge Spans	242	350	834	1,426
	# Overlapping Spans	12	8	9	29
Total	# Posts	126	123	142	391
	# Sentences	2,931	3,912	7,695	14,538
	# Tokens	55,546	78,091	98,583	232,220
	# Skill Spans	2,189	2,146	2,241	6,576
	# Knowledge Spans	807	1,418	3,828	6,053
	# Overlapping Spans	69	54	178	301
U	# Posts				126,769
	# Sentences				3,195,585
	# Tokens				460,484,670

Table 2: **Statistics of Dataset.** Indicated is the number of JPs across splits & source and their respective number of sentences, tokens, and spans. The total is reported in the cyan column and rows. We report the overall statistics of the unlabeled JPs (U) in the gray rows.

and professional and personal development.”⁴

To sum up, hard skills are usually referred to as *knowledge* components, and applying these hard skills to something is considered a *skill* component. 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.

4 SKILLSPAN Dataset

Data⁵ We continuously collected JPs via web data extraction between June 2020–September 2021. Our JPs come from the three sources:

1. **BIG:** A large job platform with various types of JPs, with several type of positions;
2. **HOUSE:** A *static* in-house dataset consisting of similar types of jobs as BIG. Dates range from 2012–2020;

⁴<http://data.europa.eu/esco/skill/A>

⁵Our data statement (Bender and Friedman, 2018) can be found in Appendix A.




Figure 2: **Violin Plots of Annotated Components.** Indicated are the distributions regarding the length of spans in each type of annotated component (i.e., length of skills and knowledge components). The white dot is the median length, the bars range from the first quartile to the third quartile, and the colored line ranges from the lower adjacent value to the higher adjacent value.

3. **TECH:** The StackOverflow JP platform that consisted mostly of technical jobs (e.g., developer positions).

We release the anonymized raw data and annotations of the parts with permissible licenses, i.e., HOUSE (from a governmental agency which is our collaborator) and TECH.⁶ 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). Table 2 shows the statistics of SKILLSPAN, with 391 annotated JPs from the three sources containing 14.5K sentences and 232.2K tokens. The unlabeled JPs (only to be released as pre-trained model) consist of 126.8K posts, 3.2M sentences, and 460.5M tokens. What stands out is that there are 2–5 times as many annotated knowledge components in TECH in contrast to the other sources, despite a similar amount of JPs. We expect this to be due the numerous KCs depicted in this domain (e.g., programming languages), while we observe considerably fewer soft skills (e.g., “work flexibly”). The amount of skills is more balanced across the three sources. Furthermore, overlapping spans follow a consistent trend among splits, with the train split containing the most.

Data Annotation We annotate competences related to SKCs in two levels as illustrated in Figure 1. We started the process in March 2021, with initial annotation rounds to construct and refine the annotation guidelines (as outlined further below).

⁶Links to our data can be found at <https://github.com/kris927b/SkillSpan>.

The annotation process spanned eight months in total. Our final annotation guidelines can be found in Appendix B. The guidelines were developed by largely following example spans given in the ESCO taxonomy. However, at this stage, we focus on span identification, and we do not take the fine-grained taxonomy codes from ESCO for labeling the spans, leaving the mapping to ESCO and taxonomy enrichment as future work.

Further Details on the Annotation Process The development of the annotation guidelines and our annotation process is depicted as follows: ① We wrote base guidelines derived from a small number of JPs. ② We had three pre-rounds consisting of three JPs each. After each round, we modified, improved and finalized the guidelines. ③ Then, we had three longer-lasting annotation rounds consisting of 30 JPs each. We re-annotated the previous 11 JPs in ① and ②. ④ After these rounds, one of the annotators (the hired linguist) annotated JPs in batches of 50. The data in ①, ②, and ③ was annotated by three annotators (101 JPs).

We used an open source text annotation tool named DOCCANO (Nakayama et al., 2018). There are around 57.5K tokens (approximately 4.6K sentences, in 101 job posts) that we calculated agreement on. The annotations were compared using Cohen’s κ (Fleiss and Cohen, 1973) between pairs of annotators, and Fleiss’ κ (Fleiss, 1971), which generalises 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 SKC, i.e., we only check the position of tag without regarding the type of the entity. The observed agreements scores over the three annotators from step ③ are between 0.70–0.75 Fleiss’ κ for both levels of calculation which is considered a *substantial agreement* (Landis and Koch, 1977) and a κ value greater than 0.81 indicates *almost perfect agreement*. Given the difficulty of this task, we consider the aforementioned κ score to be strong. Particularly, we observed a large improvement in annotation agreement from the earlier rounds (step ① and ②), where our Fleiss’ κ was 0.59 on token-level and 0.62 for the span-level.

Overall, we observe higher annotator agreement for knowledge components (3–5% higher) com-

	BIG	HOUSE	TECH
SKILL	ambitious	structured	hands-on
	proactive	teaching	communication skills
	work independently	communication skills	leadership
	attention to detail	project management	passionate
	motivated	drive	open-minded
KNOWLEDGE	full uk driving licence	english	java
	sap energy assessments	supply chain	javascript
	right to work in the uk	project management	aws
	sen	powders	docker
	acca/aca	machine learning	node.js

Table 3: **Most Frequent Skills in the Development Data.** Top-5 skill components in our data in terms of frequency on different sources. A larger example can be found in Table 8 and Table 9 (Appendix C).

pared to skills which tend to be longer. The TECH domain is the most consistent for agreement while BIG shows more variation over rounds, likely due to the broader nature of the domains of JPs.

Annotation Span Statistics A challenge of annotating spans is the length (i.e., boundary), SKCs being in different domains (e.g., business versus technical components), and frequently written differently, e.g., “being able to work together” v.s. “teamwork”). Figure 2 shows the statistics of our annotations in violin plots. For the training set, the median length (white dot) of skills is around 4 for BIG and HOUSE, for TECH this is a median of 5. In the development set, the median stays at length 4 across all sources. Another notable statistic is the upper and lower percentile of the length of skills and knowledge, indicated with the thick bars. Here, we highlight the fact that skill components could consist of many tokens, for example, up to length 7 in the HOUSE source split (see blue-colored violins). For knowledge components, the spans are usually shorter, where it is consistently below 5 tokens (see orange-colored violins). All statistics follow a similar distribution across train, development, and sources in terms of length and distribution. This gives a further strong indication that consistent annotation length has been conducted across splits and sources.

Qualitative Analysis of Annotations Qualitative differences in SKCs over the three sources are shown (lowercased) in Table 3. With respect to skill components, all sources follow a similar usage of skills. The annotated skills mostly relate to the *attitude* of a person and hence mostly consist of soft skills. With respect to knowledge components, we observe differences between the three sources. First, on the source-level, the knowledge

components vastly differ between BIG and TECH. BIG postings seem to cover more *business* related components, whereas TECH has more *engineering* components. HOUSE seems to be a mix of the other two sources. Lastly, note that both the skill and knowledge components between the splits diverge in terms of the type of annotated spans, which indicates a variation in the annotated components. We show the top-10 skills annotated in the train, development, and test splits for SKCs in Appendix C. From a syntactic perspective, skills frequently consist of noun phrases, verb phrases, or adjectives (for soft skills). Knowledge components usually consists of nouns or proper nouns, such as “python”, “java”, and so forth.

5 Experimental Setup

The task of SE is formulated as a sequence labeling problem. Formally, we consider a set of JPs \mathcal{D} , where $d \in \mathcal{D}$ is a set of sequences (i.e., entire JPs) with the i^{th} input sequence $\mathcal{X}_d^i = \{x_1, x_2, \dots, x_T\}$ and a target sequence of BIO-labels $\mathcal{Y}_d^i = \{y_1, y_2, \dots, y_T\}$ (e.g., “B-SKILL”, “I-KNOWLEDGE”, “O”). The goal is to use \mathcal{D} to train a sequence labeling algorithm $h : \mathcal{X} \mapsto \mathcal{Y}$ to accurately predict entity spans by assigning an output label y_t to each token x_t .

As baseline we consider BERT and we investigate more recent variants, and we also train models from scratch. Models are chosen due to their state-of-the-art performance, or in particular, for their strong performance on longer spans.

BERT_{base} (Devlin et al., 2019) An out-of-the-box BERT_{base} model (bert-base-cased) from the HuggingFace library (Wolf et al., 2020) functioning as a baseline.

SpanBERT (Joshi et al., 2020) A BERT-style model that focuses on span representations as opposed to single token representations. SpanBERT is trained by masking contiguous spans of tokens and optimizing two objectives: (1) masked language modeling, which predicts each masked token from its own vector representation. (2) The span boundary objective, which predicts each masked token from the representations of the unmasked tokens at the start and end of the masked span.

We train a SpanBERT_{base} model from scratch on the BooksCorpus (Zhu et al., 2015) and English Wikipedia using cased Wordpiece tokens (Wu et al., 2016). We use AdamW (Kingma and Ba, 2015) for

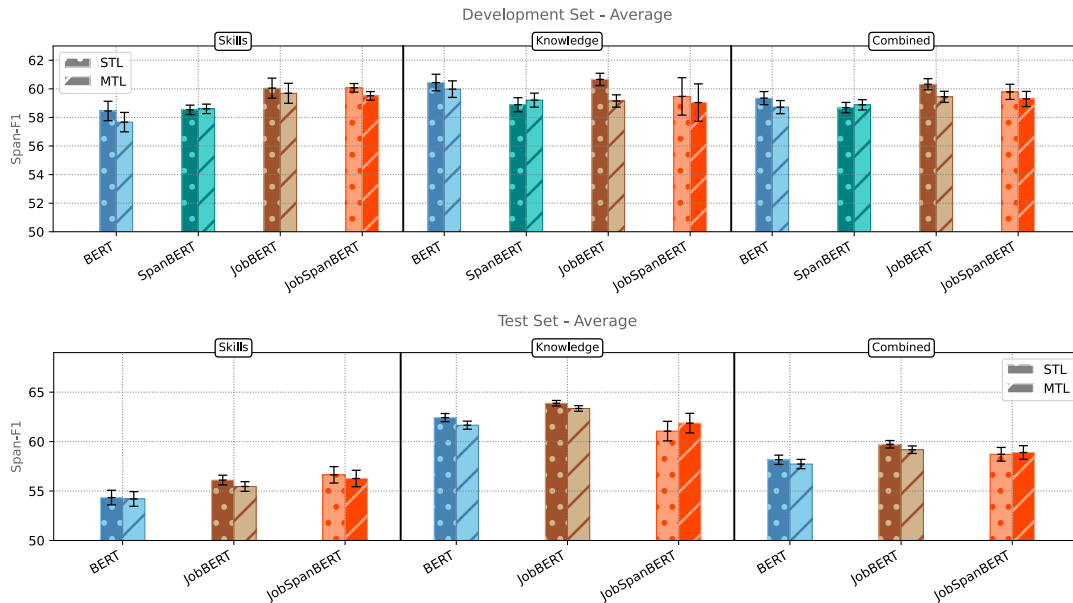


Figure 3: **Performance of Models.** We test the models on **SKILLS**, **KNOWLEDGE**, and **COMBINED**. We report the span-F1 and standard deviation (error bars) of runs on five random seeds. **Note that the y-axis starts from 50 span-F1.** STL indicates single-task learning and MTL indicates the multi-task model. Differences can be seen on the test set: JobSpanBERT performs best on SKILLS, JobBERT is best on KNOWLEDGE, and JobBERT achieves best in COMBINED. Exact numbers of the plots are in Table 5 (Appendix E).

2.4M training steps with batches of 256 sequences of length 512. The learning rate is warmed up for 10K steps to a maximum value of $1e-4$, after which it has a decoupled weight decay (Loshchilov and Hutter, 2019) of 0.1. We add a dropout rate of 0.1 across all layers. Pretraining was done on a v3-8 TPU on the GCP and took 14 days to complete. We take the official TensorFlow implementation of SpanBERT by Ram et al. (2021).

JobBERT⁷ We apply domain-adaptive pre-training (Gururangan et al., 2020) to a BERT_{base} model using the 3.2M unlabeled JP sentences (Table 2). Domain-adaptive pre-training relates to the continued self-supervised pre-training of a large language model on domain-specific text. This approach improves the modeling of text for downstream tasks within the domain. We continue training the BERT model for three epochs (default in HuggingFace) with a batch size of 16.

JobSpanBERT⁸ We apply domain-adaptive pre-training to our SpanBERT on 3.2M unlabeled JP sentences. We keep parameters identical to the vanilla SpanBERT, but change the number of steps

to 40K to have three passes over the unlabeled data.

Experiments We have 391 annotated JPs (Table 2) that we divide across three splits: Train, dev, and test set. We use 101 JPs that all three annotators annotated as the gold standard test set with aggregated annotations via majority voting. The 101 postings are divided between the sources as: 36 BIG, 33 HOUSE, and 32 TECH. The remaining 290 JPs were annotated by one annotator. We use 90 JPs (30 from each source, namely BIG, HOUSE, and TECH) as the dev. set. The remaining 200 JPs are used as the train set. The sources in the train set are divided into 60 BIG, 60 HOUSE, and 80 TECH.

Setup The data is structured as CoNLL format (Tjong Kim Sang, 2002). For the nested annotations, the skill tags are appearing only in the first column and the knowledge tags are only appearing in the second column of the file and they are allowed to overlap with each other. We perform experiments with single-task learning (STL) on either the skill or knowledge components, MTL for predicting both skill and knowledge tags at the same time, while evaluating the MTL models also on either skills or knowledge components. We used a single joint MTL model with hard-parameter sharing (Caruana, 1997). All models are with a final Conditional Random Field (CRF; Lafferty et al.,

⁷<https://huggingface.co/jjzha/jobbert-base-cased>

⁸<https://huggingface.co/jjzha/jobspanbert-base-cased>

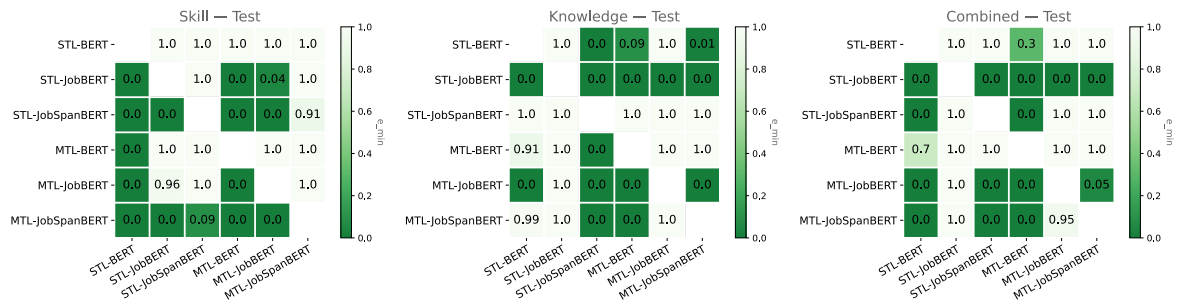


Figure 4: **Almost Stochastic Order Scores of the Test Set.** ASO scores expressed in ϵ_{\min} . The significance level $\alpha = 0.05$ is adjusted accordingly by using the Bonferroni correction (Bonferroni, 1936). Read from row to column: E.g., in COMBINED STL-JobBERT (row) is stochastically dominant over STL-BERT_{base} (column) with ϵ_{\min} of 0.00.

2001) layer. Earlier research, such as Souza et al. (2019); Jensen et al. (2021) show that BERT models with a CRF-layer improve or perform similarly to its simpler variants when comparing the overall F1 and make no tagging errors (e.g., B-tag follows I-tag). In the case of MTL we use one for each tag type (skill and knowledge). In the STL experiments we use one CRF for the given tag type.

We use the MACHAMP toolkit (van der Goot et al., 2021) for our experiments. For each setup we do five runs (i.e., five random seeds).⁹ For evaluation we use span-level precision, recall, and F1, where the F1 for the MTL setting is calculated as described in Benikova et al. (2014).

6 Results

The results of the experiments are given in Figure 3. We show the average performance of each model in F1 and respective standard deviation over the development and test split. Exact scores on each source split and other metric details are provided in Appendix E. As mentioned before, we experiment with the following settings: **SKILL**, we train and predict only on skills. **KNOWLEDGE**, train and only predict for knowledge. **COMBINED**, we merge the STL predictions of both skills and knowledge. We also train the models in an MTL setting, predicting both skills and knowledge simultaneously. We evaluate the MTL model on both **SKILL** and **KNOWLEDGE** separately, and also compare it against the aggregated STL predictions.

Performance on Development Set In Figure 3, we show the results on the development set in the upper plot. We observe similar performance between the domain-adapted STL models—JobBERT and JobSpanBERT—have similar span-F1 for

⁹For reproducibility, we refer to Appendix D.

SKILL: 60.05 ± 0.70 vs. 60.07 ± 0.70 . In contrast, for **KNOWLEDGE**, BERT_{base} and JobBERT are closest in predictive performance: 60.44 ± 0.58 vs. 60.66 ± 0.43 . In the **COMBINED** setting, JobBERT performs highest with a span-F1 of 60.32 ± 0.39 . On average, JobBERT performs best over all three settings. Surprisingly, the models for both **SKILL** and **KNOWLEDGE** perform similarly (around 60 span-F1), despite the sources’ differences in properties and length Figure 2. In addition, we find that MTL is not performing better than STL across sources. For exact numbers and source-level (i.e., BIG, HOUSE, TECH), we refer to Appendix E.

Performance on Test Set We select the best performing models in the development set evaluation and apply it to the test set. Results are in Figure 3 in the bottom plot. Since JobBERT and JobSpanBERT are performing similarly, we apply both to the test set and BERT_{base}. We observe a deviation from the development set to the test set: JobSpanBERT $60.07 \pm 0.30 \rightarrow 56.64 \pm 0.83$ on **SKILL**, JobBERT $60.66 \pm 0.43 \rightarrow 63.88 \pm 0.28$ on **KNOWLEDGE**. For **COMBINED**, JobBERT performs slightly worse: $60.32 \pm 0.39 \rightarrow 59.73 \pm 0.38$. Similar to the development set, we find that on all three methods of evaluation (i.e., **SKILL**, **KNOWLEDGE**, and **COMBINED**), STL still outperforms MTL. For **SKILL** and **KNOWLEDGE**, STL is almost stochastically dominant over MTL (i.e., significant), and for **COMBINED** there is stochastic dominance of STL over MTL, indicated in the next paragraph.

Significance We compare all pairs of models based on five random seeds each using Almost Stochastic Order (ASO; Dror et al., 2019) tests with a confidence level of $\alpha = 0.05$. The ASO scores of the test set are indicated in Figure 4. We show that MTL-JobSpanBERT for **SKILL** shows

almost stochastic dominance ($\epsilon_{\min} < 0.5$) over all other models. For KNOWLEDGE and COMBINED, We show that STL-JobBERT is stochastically dominant ($\epsilon_{\min} = 0.0$) over *all* the other models. For more details, we refer to [Appendix F](#) for ASO scores on the development set.

7 Discussion

What Did Not Work Additionally, we experiment whether representing the entire JP for extracting tokens yields better results than the experiments so far, which were sentence-by-sentence processing setups. To handle entire JPs and hence much longer sequences we use a pre-trained Longformer_{base} (Beltagy et al., 2020) model. The document length we use in the experiments is 4096 tokens. Results of the Longformer on the test set are lower: For skills, JobSpanBERT against Longformer results in 56.64 ± 0.83 vs. 52.55 ± 2.39 . For KNOWLEDGE, JobBERT against Longformer shows 63.88 ± 0.28 vs. 57.26 ± 1.05 . Last, for COMBINED, JobBERT against Longformer results in 59.73 ± 0.38 vs. 55.05 ± 0.71 . This drop in performance is difficult to attribute to a concrete reason: e.g., the Longformer is trained on more varied sources than BERT, but not specifically for JPs, which may have contributed to this gap. Since the vanilla Longformer already performs worse than BERT_{base} overall, we did not opt to apply domain-adaptive pre-training. Overall, we show that representing the full JP is not beneficial for SE, at least not in the Longformer setup tested here.

Continuous Pretraining helps SE As previously mentioned, due to the domain specialization of the domain-adapted pre-trained BERT models, they predict more skills and frequently perform better in terms of precision, recall, and F1 as compared to their non-adaptive counterparts. This is especially encouraging as we confirm findings that continuous pre-training helps to adapt models to a specific domain (Alsentzer et al., 2019; Lee et al., 2020; Gururangan et al., 2020; Nguyen et al., 2020). However, there are exceptions. Particularly in [Table 5](#) on TEST for KNOWLEDGE, BERT_{base} comes closer in predictive performance to JobBERT (difference of 1.5 F1) than on SKILLS. Our intuition is that knowledge components are often already in the pre-training data (e.g., Wikipedia pages of certain competences like Python, Java etc.) and therefore adaptive pre-training does not substantially boost performance.

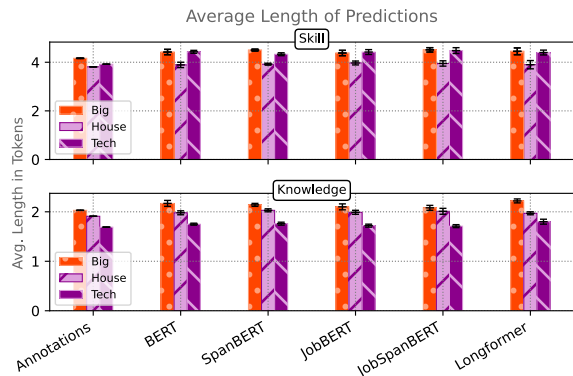


Figure 5: **Average Length of Predictions of Single Models.** We show the average length of the predictions versus the length of our annotated skills and knowledge components on the *test set* and the total number of predicted skills and knowledge tags in each respective split (#). There is a consistent trend over the three sources.

Difference in Length of Predictions The main motivation of selecting models optimized for long spans was the length of the annotations ([Figure 2](#)). We investigate the average length of predictions of each model ([Figure 5](#)) to find out whether the models that are adapted to handle longer sequences truly predict longer spans. Interestingly, the average length of predicted skills are longer than the annotations over all three sources. There is a consistent trend among SKILL: BIG and TECH have similar length over predictions (>4), while HOUSE is usually lower than length 3. For both BIG and TECH, JobSpanBERT predicts the longest skill spans (4.51 and 4.48 respectively). We suspect due to the domain-adaptive pre-training on JPs, it improved the span prediction performance. In contrast, the Longformer predicts shorter spans. Note that the Longformer is not domain-adapted to JPs.

Regarding KNOWLEDGE, there is also a consistent trend: BIG has the overall longest prediction length while TECH has the lowest. The Longformer predicts the longest spans on average for BIG and TECH. Knowledge components are representative of a normal-length NER task and might not need a specialized model for long sequences. We show the exact numbers in [Table 7](#) ([Appendix E](#)) and the number of predicted SKILL and KNOWLEDGE: JobBERT and JobSpanBERT have higher recall than the other models.

Performance per Span Length SKILLS are generally longer than KNOWLEDGE components in our dataset ([Figure 2](#)). The previous overall results on the test set ([Figure 3](#)) show that performance on




Figure 6: **Average Span-F1 per Span Length.** We bucket the performance of JobBERT according to the length of the spans until 10 tokens and show the performance on each length, averaged over five random seeds. Indicated per bar is the support. The model performs best on medium-length skill spans (i.e., spans with token length of 4-5). For knowledge spans, on average, it performs best on short-length spans (i.e., spans with token length of 1-2).

SKILL is substantially lower than KNOWLEDGE. We therefore investigate whether this performance difference is attributed to the longer spans in SKILL. In Figure 6, we show the average performance of the best performing model (JobBERT) on the three sources (test set) based on the gold span length, until a length of 10.

In SKILL components (upper plot), we see much support for spans with length 1 and 2, which then lowers once the spans become longer. Spans with length of 1 shows low performance on BIG and TECH (around 40 span-F1), which influences the total span-F1. Short skills are usually soft skills, such as “passionate”, which can be used as a skill or not. This might confuse the model. In contrast, performance effectively stays similar (around 60 span-F1) for span length of 2 till 7 for all sources. Afterwards, it drops in performance. Thus, the weak performance on SKILL seem to be due to lower performance on the short spans.

For the KNOWLEDGE components (lower plot), they are generally shorter. We see that there is a gap in support between the sources, TECH has a larger number of gold labels compared to BIG and HOUSE. Unlike soft skills, KCs usually consist of proper nouns such as “Python”, “Java”, and so forth, which connects to the high performance on TECH (around 76 span-F1). Furthermore, support for spans longer than 2 drops considerably. In this

case, if the model predicts a couple of instances correctly, it would substantially increase span-F1. Contrary to SKILL, high performance of KNOWLEDGE can be attributed to its strong performance on short spans.

8 Conclusion

We present a novel dataset for skill extraction on English job postings—SKILLSPAN—and domain-adapted BERT models—JobBERT and JobSpanBERT. We outline the dataset and annotation guidelines, created for hard *and* soft skills annotation on the *span-level*. Our analysis shows that domain-adaptive pre-training helps to improve performance on the task for both skills and knowledge components. Our domain-adapted JobSpanBERT performs best on skills and JobBERT on knowledge. Both models achieve almost stochastic dominance over all other models for skills and knowledge extraction, whereas JobBERT in the STL setting achieves stochastic dominance over other models.

With the rapid emergence of new competences, our new approach of skill extraction has future potential, e.g., to enrich knowledge bases such as ESCO with unseen skills or knowledge components, and in general, contribute to providing insights into labor market dynamics. We hope our dataset encourages research into this emerging area of computational job market analysis.

Acknowledgements

We thank Google’s TFRC for their support in providing TPUs for this research. Furthermore, 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 (DFR) grant 9131-00019B.

References

- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Krisztian Balog, Yi Fang, Maarten De Rijke, Pavel Serdyukov, and Luo Si. 2012. Expertise retrieval. *Foundations and Trends in Information Retrieval*, 6(2–3):127–256.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *ArXiv preprint*, abs/2004.05150.
- Emily M. Bender and Batya Friedman. 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.
- Darina Benikova, Chris Biemann, and Marc Reznicek. 2014. [NoSta-D named entity annotation for German: Guidelines and dataset](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 2524–2531, Reykjavik, Iceland. European Language Resources Association (ELRA).
- 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 *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5832–5842, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Carlo Bonferroni. 1936. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62.
- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Mariia Chernova. 2020. [Occupational skills extraction with FinBERT](#). *Master’s Thesis*.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rotem Dror, Segev Shlomov, and Roi Reichart. 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.
- Peter Elias. 1997. Occupational classification (isco-88): Concepts, methods, reliability, validity and cross-national comparability. Technical report, OECD Publishing.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Akshay Gugnani and Hemant Misra. 2020. [Implicit skills extraction using document embedding and its use in job recommendation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 13286–13293. AAAI Press.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Xiaochuang Han and Jacob Eisenstein. 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.

- Sepp Hochreiter, Jürgen Schmidhuber, and Corso Elvezia. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Faizan Javed, Phuong Hoang, Thomas Mahoney, and Matt McNair. 2017. [Large-scale occupational skills normalization for online recruitment](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4627–4634. AAAI Press.
- Kristian Nørgaard Jensen, Mike Zhang, and Barbara Plank. 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.
- Shanshan Jia, Xiaoan Liu, Ping Zhao, Chang Liu, Lianyong Sun, and Tao Peng. 2018. Representation of job-skill in artificial intelligence with knowledge graph analysis. In *2018 IEEE Symposium on Product Compliance Engineering-Asia (ISPC-Asia)*, pages 1–6. IEEE.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Imane Khaouja, Ismail Kassou, and Mounir Ghogho. 2021. A survey on skill identification from online job ads. *IEEE Access*, 9:118134–118153.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Ilkka Kivimäki, Alexander Panchenko, Adrien Dessy, Dries Verdegem, Pascal Francq, Hugues Bersini, and Marco Saerens. 2013. [A graph-based approach to skill extraction from text](#). In *Proceedings of TextGraphs-8 Graph-based Methods for Natural Language Processing*, pages 79–87, Seattle, Washington, USA. Association for Computational Linguistics.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001), Williams College, Williamstown, MA, USA, June 28 - July 1, 2001*, pages 282–289. Morgan Kaufmann.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Martin le Vrang, Agis Papantoniou, Erika Pauwels, Pieter Fannes, Dominique Vandestein, and Johan De Smedt. 2014. Esco: Boosting job matching in europe with semantic interoperability. *Computer*, 47(10):57–64.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Shan Li, Baoxu Shi, Jaewon Yang, Ji Yan, Shuai Wang, Fei Chen, and Qi He. 2020. [Deep job understanding at linkedin](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 2145–2148. ACM.
- Liting Liu, Wenzheng Zhang, Jie Liu, Wenxuan Shi, and Yalou Huang. 2021. Learning multi-graph neural network for data-driven job skill prediction. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang. 2018. [doccano: Text annotation tool for human](#). Software available from <https://github.com/doccano/doccano>.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Ori Ram, Yuval Kirstain, Jonathan Berant, Amir Globerson, and Omer Levy. 2021. [Few-shot question answering by pretraining span selection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3066–3079, Online. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2018. [Why comparing single performance scores does not allow to draw conclusions about machine learning approaches](#). *ArXiv preprint*, abs/1803.09578.
- Luiza Sayfullina, Eric Malmi, and Juho Kannala. 2018. Learning representations for soft skill matching. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 141–152.
- Baoxu Shi, Jaewon Yang, Feng Guo, and Qi He. 2020. [Salience and market-aware skill extraction for job targeting](#). In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 2871–2879. ACM.

- Ellery Smith, Martin Braschler, Andreas Weiler, and Thomas Habertuer. 2019. Syntax-based skill extractor for job advertisements. In *2019 6th Swiss Conference on Data Science (SDS)*, pages 80–81. IEEE.
- Ellery Smith, Andreas Weiler, and Martin Braschler. 2021. Skill extraction for domain-specific text retrieval in a job-matching platform. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 116–128. Springer.
- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2019. Portuguese named entity recognition using bert-crf. *ArXiv preprint*, abs/1909.10649.
- Damian A Tamburri, Willem-Jan Van Den Heuvel, and Martin Garriga. 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.
- Erik F. Tjong Kim Sang. 2002. Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition. In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Dennis Ulmer. 2021. deep-significance: Easy and Better Significance Testing for Deep Neural Networks. <https://github.com/Kaleidophon/deep-significance>.
- Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 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.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *ArXiv preprint*, abs/1609.08144.
- Meng Zhao, Faizan Javed, Ferosh Jacob, and Matt McNair. 2015. SKILL: A system for skill identification and normalization. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 4012–4018. AAAI Press.
- Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.

A Data Statement SKILLSPAN

Following [Bender and Friedman \(2018\)](#), the following outlines the data statement for SKILLSPAN:

- A. **CURATION RATIONALE:** Collection of job postings in the English language for span-level sequence labeling, to study the impact of sequence labeling on the extraction of skill and knowledge components 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) and British (en-GB) English 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 or British English 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 job posting data in TECH is from Stackoverflow jobs, and is licensed under the CC BY-SA license. The job posting data from HOUSE is from our collaborators: The Danish Agency for Labour Market and Recruitment (STAR).

B Annotation Guidelines

B.1 Span Specifications

Legend: Skill, Knowledge, “•” indicates an example sentence.
1. A skill starts with a VERB , otherwise (ADJECTIVE) + NOUN
1.1 Modal verbs are not tagged: <ul style="list-style-type: none">• Can [put personal touch on the menu]_{SKILL} .• Will [train new staff]_{SKILL} .
2. Split up phrases with prepositions and/or conjunctions
2.1 Unless the conjunction coordinates two nouns functioning as one argument: <ul style="list-style-type: none">• [Coordinate parties and conferences]_{SKILL} .
2.2 Do not tag skills with anaphoric pronouns, only tag preceding skill: <ul style="list-style-type: none">• [Prioritizing tasks]_{SKILL} and identifying those that are most important.
2.3 Split nouns and adjectives that are coordinated if they do not have a verb attached: <ul style="list-style-type: none">• Be [inquisitive]_{SKILL} and [proactive]_{SKILL} .• Prior in-house experience with [media]_{KNOWLEDGE} , [publishing]_{KNOWLEDGE} or [internet companies]_{KNOWLEDGE} .
2.4 If there is a listing of skill tags and they lead up to different subtasks, we split them: <ul style="list-style-type: none">• [keep up the high level of quality in our team]_{SKILL} through [reviews]_{SKILL} , [pairing]_{SKILL} and [mentoring]_{SKILL} .
3. If there is relevant information appended after irrelevant information (e.g., info specific to a company) we try to make the skill as short as possible: <ul style="list-style-type: none">• [providing the best solution]_{SKILL} for Siemens Gamesa in a very [structured]_{SKILL} and [analytic]_{SKILL} manner.
4. Note also the words skills and knowledge can be included in the span of the component if leaving it out makes it nonsensical: <ul style="list-style-type: none">• [personal skills]_{SKILL} → just [personal] would make it nonsensical.
5. Parentheses after a skill tag are included if they elaborate the component before them or if they are an abbreviation of the component.

6. Inclusion of adverbials in components. Adverbials are included if it concerns the manner of doing something. All others are excluded:

- ~~like to~~ [solve technical challenges independently]_{SKILL} .
- [communicates openly]_{SKILL} .
- [striving for the best]_{SKILL} ~~in all that they do.~~
- [Deliver first class customer service]_{SKILL} ~~to our guests.~~
- [Making the right decisions]_{SKILL} ~~early in the process.~~

7. Attitudes as skills. We annotate attitudes as a skill:

- a [can-do-approach]_{SKILL} → we leave out articles from the attitude.

8. Attitudes are not tagged if they contain skill/knowledge components—then only the span of the skill is tagged.

- ~~like to~~ [solve technical challenges independently]_{SKILL} .
- ~~Passion for~~ [automation]_{KNOWLEDGE} .
- ~~enjoy~~ [working in a team]_{SKILL} .

9. Miscellaneous:

9.1 Do not tag ironic skills (e.g., lazy).

9.2 Avoid nesting of skills, annotate it as one span.

9.3 We annotate all skills that are part of sections such as “requirements”, “good-to-haves”, “great-to-knows”, “optionals”, “after this x months of training you’ll be able to...”, “At the job you’re going to...”.

9.4 When there is a general standard that can be added to the skill, we add these:

- [Process payments according to the [...] standards]_{SKILL} .

B.2 Knowledge Specifications

1. **Rule-of-thumb:** knowledge is something that one possesses, and cannot (usually) physically execute:

- [Python]_{KNOWLEDGE} (programming language).
- [Business]_{KNOWLEDGE} .
- [Relational Databases]_{KNOWLEDGE} .

2. If there is a component between parentheses that belongs to the knowledge component, we add it:

- [(non-) relational databases]_{KNOWLEDGE} .
- [Driver License (UK/EU)]_{KNOWLEDGE} .

3. **Licenses and certifications:** We add the additional words “certificate”, “card”, “license”, et cetera. to the knowledge component.

4. If the knowledge component looks like a skill, but the preceding verb is vague and empty (e.g., *follow*, *use*, *comply with*, *work with*) → only tag the knowledge component:

- Comply with [Food Code of Practice]_{KNOWLEDGE} .
- Work with [AWS infrastructure]_{KNOWLEDGE} .

5. We annotate only specified knowledge components:

- [MongoDB]_{KNOWLEDGE} or other [NoSQL database]_{KNOWLEDGE} .
- [JEST]_{KNOWLEDGE} or other test libraries. → “other test libraries” is under-specified.

6. Knowledge components can be nested in skill components.

- [Design, execution and analysis of [phosphoproteomics]_{KNOWLEDGE} experiments]_{SKILL} .

7. If all components coordinate/share one knowledge tag, we annotate it as one:

- [application, data and infrastructure architecture]_{KNOWLEDGE} . → The knowledge tags coordinate to “architecture”.
- [chemical/biochemical engineering]_{KNOWLEDGE} .

8. If there is a listing of knowledge tags, we annotate all knowledge tags separately:

- [Bachelor Degree]_{KNOWLEDGE} in [Mathematics]_{KNOWLEDGE} , [Computer Science]_{KNOWLEDGE} , or [Engineering]_{KNOWLEDGE} .

B.3 Other Specifications

1. Rule-of-thumb: If in doubt, annotate it as a skill.
2. We are preferring skills over knowledge components.
3. We prioritize skills over attitudes; if there is a skill within the attitude, only tag the skill: <ul style="list-style-type: none"> • Passionate around [solving business problems]_{SKILL} through [innovation & engineering practices]_{KNOWLEDGE}.
4. Skill or knowledge components in the top headlines of the JP are not tagged (e.g., title of a JP). If it is a sub-headline or in the rest of the posting, tag it.
5. We try to keep the skill/knowledge components as short as possible (i.e., exclude information at the end if it makes it too specific for the job).
6. We do not include “fluff” and “triggers” (i.e., words that indicate a skill or knowledge component will follow: “advanced knowledge of [...]KNOWLEDGE”) around the components, including degree. This goes for both before and after: <ul style="list-style-type: none"> • Working proficiency in [developmental toolsets]_{KNOWLEDGE}. • Advanced knowledge of [application data and architecture infrastructure]_{KNOWLEDGE} disciplines. • [Manual handling]_{SKILL} tasks. • [CI/CD]_{KNOWLEDGE} experience. • You master [English]_{KNOWLEDGE} on level C1. • Proficient in [Python]_{KNOWLEDGE} and [English]_{KNOWLEDGE}. • Fluent in spoken and written [English]_{KNOWLEDGE}.
7. Pay attention to expressions such as “participation in...”, “contributing”, and “transfer (knowledge)”. These are usually not considered skills. <ul style="list-style-type: none"> • Contribute to the enjoyable and collaborative work environment. • Participation in the Department’s regular research activities. • Desire to be part of something meaningful and innovative.
8. Skills and Knowledge components that are found in not-so-straightforward places (e.g., project descriptions) are annotated as well, if they relate to the position.
9. In the pattern of “skill” followed by some elaboration, see if it can be annotated with a skill and a knowledge tag: <ul style="list-style-type: none"> • [Ensure food storage and preparation areas are maintained]_{SKILL} according to [Health & Safety and Audit standards]_{KNOWLEDGE}.

10. Occupations and positions in companies/academia should be excluded.
11. If there's a knowledge/skill component in the position, we exclude it as well. <ul style="list-style-type: none">• Experienced Java Engineer. → completely untagged.
12. Only annotate the skills that are <u>related</u> to the position. 12.1. This includes skills that are specific for the position as well (e.g., skills of a ruminants professor versus math professor). 12.2 Also skills that the person for the position is expected to do in the future. 12.3 This does <u>not</u> include skills, knowledge or attitudes describing only the company, the group you will join in the department, and so on. <u>Only annotate</u> if it is specified or implied that the employee should possess the skill as well.
13. We annotate industries and fields (that the employee will be working in) as knowledge components.

C Type of Skills Annotated

In both Table 8 and Table 9, we show the top-10 skill and knowledge components that have been annotated. We split the top-10 among the data splits (i.e., train, development, and test set), and also between source splits (i.e., BIG, HOUSE, TECH).

D Reproducibility

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 4: Hyperparameters of MACHAMP.

We use the default hyperparameters in MACHAMP (van der Goot et al., 2021) as shown in Table 4. 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[®] TITAN X (Pascal) 12 GB GPU and an Intel[®] Xeon[®] Silver 4214 CPU.

E Exact Number of Performance

In Table 5, we show the exact numbers of the plot indicated in Figure 3. In addition, we also show the results of each respective split.

For the STL models, we observe differences in performances over the sources which is particularly pronounced for knowledge components: The TECH source is the easiest to process (and has most SKCs), while SKCs identification performance is the lowest for BIG. This might be due to the broad nature of this source.

In the exact results table (Table 5) we add a (†) next to the highest span-F1 if the model is truly stochastically dominant ($\epsilon_{\min} = 0.0$) over *all* the other models. (*) denotes that the best model achieved *almost stochastic dominance* ($\epsilon_{\min} < 0.5$) over—at minimum—one other model (e.g., in TEST rows w.r.t COMBINED: MTL-JobBERT \succeq MTL-JobSpanBERT with $\epsilon_{\min} = 0.06$) and stochastically dominant over the rest.

In Table 6, we report the precision and recall of the models, SKILL and KNOWLEDGE show the precision and recall of the STL models. MULTI shows the precision and recall of the MTL models.

Last, in Table 7, we show the exact numbers of the length of predictions Figure 5. We also add the number of predicted SKILL and KNOWLEDGE Overall, JobBERT and JobSpanBERT predict more skills in general than the other models. This is also the case for knowledge components. We hypothesize that this might be due to the BERT models now being more specialized towards the JP domain and recognizing more SKCs.

F Significance Testing

Recently, the ASO test (Dror et al., 2019)¹⁰ 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 took 0.05 for the predefined significance level. In Figure 7, we show the ASO scores on the development set.

¹⁰Implementation of Dror et al. (2019) can be found at <https://github.com/Kaleidophon/deep-significance> (Ulmer, 2021)

	Evaluation →	SKILL		KNOWLEDGE		COMBINED	
Src.	↓ Model, Task →	STL	MTL	STL	MTL	STL (*2)	MTL
BIG	BERT _{base}	59.55±0.97	58.88±1.14	50.68±3.25	51.10±1.67	57.46±1.19	57.00±0.91
	SpanBERT	59.78±0.44	60.02±2.15	50.65±2.32	51.79±2.12	57.71±0.53	58.00±2.07
	JobBERT	60.60±0.81	59.76±0.60	50.29±1.86	47.59±1.11	58.19±0.49	56.75±0.50
	JobSpanBERT	60.16±0.61	59.44±1.11	45.20±2.76	47.69±3.38	56.56±0.49	56.58±0.63
HOUSE	BERT _{base}	56.83±1.29	55.89±1.90	55.00±1.11	54.05±1.00	56.17±0.92	55.20±1.35
	SpanBERT	57.54±1.08	57.30±0.84	52.01±1.72	51.48±1.01	55.55±1.10	55.09±0.74
	JobBERT	59.81±1.17	59.97±0.85	54.94±1.15	54.23±2.60	58.02±0.93	57.80±1.50
	JobSpanBERT	59.97±1.03	59.62±0.74	55.66±1.51	53.10±1.27	58.37±1.07	57.14±0.56
TECH	BERT _{base}	59.05±0.71	58.34±0.75	64.08±1.04	63.77±1.18	62.10±0.67	61.65±0.62
	SpanBERT	58.39±0.46	58.61±1.14	62.68±0.60	63.40±0.93	61.02±0.35	61.56±0.81
	JobBERT	59.81±0.75	59.36±0.90	64.57±0.42	63.15±0.94	62.69±0.40	61.67±0.90
	JobSpanBERT	60.09±1.43	59.48±0.61	63.40±1.51	63.23±0.64	62.09±0.85	61.80±0.54
AVERAGE	BERT _{base}	58.45±0.68	57.67±1.01	60.44±0.58	59.98±0.75	59.35±0.46	58.72±0.48
	SpanBERT	58.53±0.33	58.60±0.83	58.89±0.49	59.21±0.78	58.69±0.36	58.88±0.64
	JobBERT	60.05±0.70	59.69±0.62	60.66±0.43*	59.15±1.07	60.32±0.39*	59.44±0.81
	JobSpanBERT	60.07±0.30†	59.51±0.68	59.47±1.31	59.04±0.65	59.79±0.53	59.29±0.43
TEST	BERT _{base}	54.34±0.74	54.20±0.68	62.43±0.41	61.66±0.83	58.16±0.47	57.73±0.66
	JobBERT	56.11±0.49	55.46±0.75	63.88±0.28*	63.35±0.30	59.73±0.38†	59.18±0.37
	JobSpanBERT	56.64±0.83*	56.27±0.55	61.06±0.99	61.87±0.55	58.72±0.69	58.90±0.48

Table 5: **Performance of Models.** We test the models on **skills**, **KNOWLEDGE**, and **COMBINED** (MTL). We report the span-F1 and standard deviation of runs on five random seeds on the *development set* (**AVERAGE**, in gray). Results on the *test set* are below in the **TEST** rows (in cyan). **STL** indicates single-task learning and **MTL** indicates the multi-task model. **Bold** numbers indicate best performing model in that experiment. A (†) means that it is stochastically dominant over *all* the other models. (*) denotes *almost stochastic dominance* ($\epsilon_{\min} < 0.5$) over—at minimum—one other model.

	Evaluation →	SKILL		KNOWLEDGE		MULTI	
Src.	↓ Model	Precision	Recall	Precision	Recall	Precision	Recall
BIG	BERT _{base}	57.09±1.70	62.27±1.28	43.95±4.17	60.00±1.65	52.63±1.32	62.19±0.87
	SpanBERT	58.28±0.59	61.36±0.68	45.80±2.89	56.82±3.39	54.02±1.81	62.63±2.60
	JobBERT	57.90±1.25	63.59±0.99	43.45±1.98	59.84±3.44	51.13±0.48	63.74±0.79
	JobSpanBERT	58.39±1.03	62.09±1.85	38.55±3.12	54.76±3.18	52.22±0.35	61.75±1.22
HOUSE	BERT _{base}	55.95±2.46	57.79±0.67	52.84±0.65	57.42±2.76	51.65±1.11	59.28±2.07
	SpanBERT	56.70±1.59	58.44±1.16	49.87±2.57	54.49±3.09	52.27±0.64	58.25±1.50
	JobBERT	58.16±1.30	61.56±1.53	51.18±2.18	59.37±1.34	53.72±1.57	62.56±1.47
	JobSpanBERT	59.04±0.85	60.99±2.58	51.36±2.70	60.84±1.19	53.91±0.77	60.79±0.54
TECH	BERT _{base}	58.28±1.30	59.89±1.39	60.79±1.89	67.79±1.20	58.19±1.12	65.55±0.75
	SpanBERT	58.62±0.32	58.16±0.76	59.43±1.21	66.35±1.18	58.34±0.97	65.17±1.41
	JobBERT	58.81±1.38	60.88±1.51	61.38±1.11	68.14±1.36	57.69±0.93	66.25±0.90
	JobSpanBERT	59.86±3.07	60.40±0.68	59.78±2.43	67.57±1.97	58.26±0.82	65.82±0.91
AVERAGE	BERT _{base}	57.11±1.65	59.90±0.95	56.86±1.33	64.54±1.31	55.02±0.85	62.98±0.93
	SpanBERT	57.85±0.65	59.23±0.52	55.65±1.09	62.58±1.56	55.61±0.61	62.58±1.25
	JobBERT	58.29±1.08	61.94±1.16	56.73±1.41	65.22±1.03	55.03±0.84	64.62±0.77
	JobSpanBERT	59.11±1.59	61.12±1.49	55.11±2.41	64.66±1.38	55.64±0.56	63.46±0.69
TEST	BERT _{base}	56.02±1.50	52.79±1.18	59.09±0.85	66.20±1.69	55.82±1.03	59.79±0.87
	JobBERT	55.94±1.19	56.29±0.49	60.03±1.13	68.30±1.46	55.87±0.29	62.89±0.56
	JobSpanBERT	57.57±1.24	55.77±1.65	57.83±1.03	64.71±2.10	57.06±0.74	60.89±0.42

Table 6: **Precision and Recall of Models.** We test the models on *skills*, *knowledge*, and *multi-task* setting. We report the average precision, recall and standard deviation of runs on five random seeds on the *development set* (**AVERAGE**). Results on the *test set* are below in the **TEST** rows.

Source →	BIG		HOUSE		TECH	
↓ Model	SKILLS (#)	KNOWLEDGE (#)	SKILLS (#)	KNOWLEDGE (#)	SKILLS (#)	KNOWLEDGE (#)
ANNOTATIONS	4.16 (634)	2.03 (242)	3.81 (637)	1.91 (350)	3.92 (459)	1.69 (834)
BERT _{base}	4.42±0.11 (628)	2.17±0.06 (307)	3.89±0.11 (615)	1.98±0.04 (461)	4.43±0.06 (449)	1.75±0.02 (885)
SpanBERT	4.50±0.04 (621)	2.14±0.03 (298)	3.92±0.04 (597)	2.03±0.03 (441)	4.33±0.06 (444)	1.76±0.03 (869)
JobBERT	4.38±0.11 (670)	2.10±0.06 (313)	3.97±0.08 (650)	1.99±0.04 (470)	4.42±0.10 (479)	1.72±0.03 (932)
JobSpanBERT	4.51±0.09 (629)	2.08±0.05 (313)	3.95±0.11 (623)	2.01±0.06 (452)	4.48±0.12 (439)	1.71±0.03 (875)
Longformer	4.45±0.14 (653)	2.22±0.04 (298)	3.90±0.17 (639)	1.97±0.03 (483)	4.40±0.10 (472)	1.80±0.05 (864)

Table 7: **Average Length of Predictions of Single Models.** We show the average length of the predictions versus the length of our annotated skills and knowledge components on the *test set* and the total number of predicted skills and knowledge tags in each respective split (#).




Figure 7: **Almost Stochastic Order Scores of the Development Set.** ASO scores expressed in ϵ_{\min} . The significance level $\alpha = 0.05$ is adjusted accordingly by using the Bonferroni correction (Bonferroni, 1936). Read from row to column: E.g., STL-JobBERT (row) is stochastically dominant over STL-BERT_{base} (column) with ϵ_{\min} of 0.00.

Skills			
Src.	Train	Test	
Big	<p>enthusiastic flexible team player friendly attention to detail communicator passionate communication confident flexible approach</p>	<p>ambitious proactive work independently attention to detail motivated reliable flexible willingness to learn self-motivated work as part of an established team</p>	<p>customer service communicator flexible attention to detail ambitious design and refine every touchpoint of the customer journey enable inclusion communicate effectively interpersonal skills proactive</p>
HOUSE	<p>communication skills motivated structured proactive analytical communication self-driven team player teaching curious</p>	<p>structured teaching communication skills project management drive problem solving communication visit customers curious work independently</p>	<p>teaching research communication skills outgoing flexible energetic responsible enthusiastic team player communication</p>
TECH	<p>communication skills passionate apply your depth of knowledge and expertise partner continuously with your many stakeholders solving business problems through innovation and engineering practices work in large collaborative teams hands-on building an innovative culture team player develop</p>	<p>hands-on communication skills leadership passionate open-minded code reviews independent software development pioneer new approaches analytical skills</p>	<p>solving business problems apply your depth of knowledge and expertise partner continuously with your many stakeholders achieve organizational goals building an innovative culture stay focused on common goals work in large collaborative teams design development communicate</p>

Table 8: Most Frequent Skills in the Data. Top-10 skill components in our data in terms of frequency.

	Knowledge		
Src.	Train	Development	Test
BIG	english driving license excel cscs card maths ppc service design uk/emea policies bachelor's degree computer science	full uk driving licence sap energy assessments right to work in the uk sen acca/aca professional kitchen cra calculations email marketing qualitative and quantitative social research methods care setting	strategic planning english cscs card pms reservation systems keynote illustrator aba sen full driving license
HOUSE	english engineering computer science product management python finance project management agile danish javascript	english supply chain project management powders machine learning phd degree muscle models with learning and adaptation walking robots model rules capacity development	english danish business java marketing plm production supply chain economics excel
TECH	javascript python java agile financial services node.js english kubernetes cloud docker	java javascript aws docker node.js typescript react linux amazon-web-services devops	java python .net financial services c# javascript cloud english reactjs automation

Table 9: **Most Frequent Knowledge in the Data.** Top-10 knowledge components in our data in terms of frequency.