Computational Job Market Analysis

with Natural Language Processing

Mike Zhang, PhD Defence, 7 March 2024

IT UNIVERSITY OF CPH

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

Max™


COOKIES
CLASSIC · WAFERS · BRITTLE
EAT LATER
BREAKFAST
NO LINES. NO CHECKOUT.
(NO, SERIOUSLY.)
JUST
WALK
OUT
SHOPPING
amazon go
EXIT
Thanks for shopping. You're


JORDI
ARENAS
I CLAVELL
(MATARÓ, 1920 - 1998)
TAMARIU, S/D
OLI SOBRE TELA
DESCOBREIX MÉS
SOBRE L'AUTOR
tant en essència com per la
immensitat de la natura com
un violent onatge que trenca
nificència es tradueix en un
ina, on també hi té presència
persones accentua aquesta
l qual ens n'allunyem quan
de color, llum i moviment.


AA-B77 SP1
AA-B13 SP1


SCHOOL
SPEED LIMIT 15
SCHOOL
WAYMO
JAGUAR
TCP 38152A
8MPK559


EXIT
Thanks for shopping. You're
COOKIES
EAT LATER
BREAKFAST
JUST
WALK
OUT
SHOPPING
amazon go

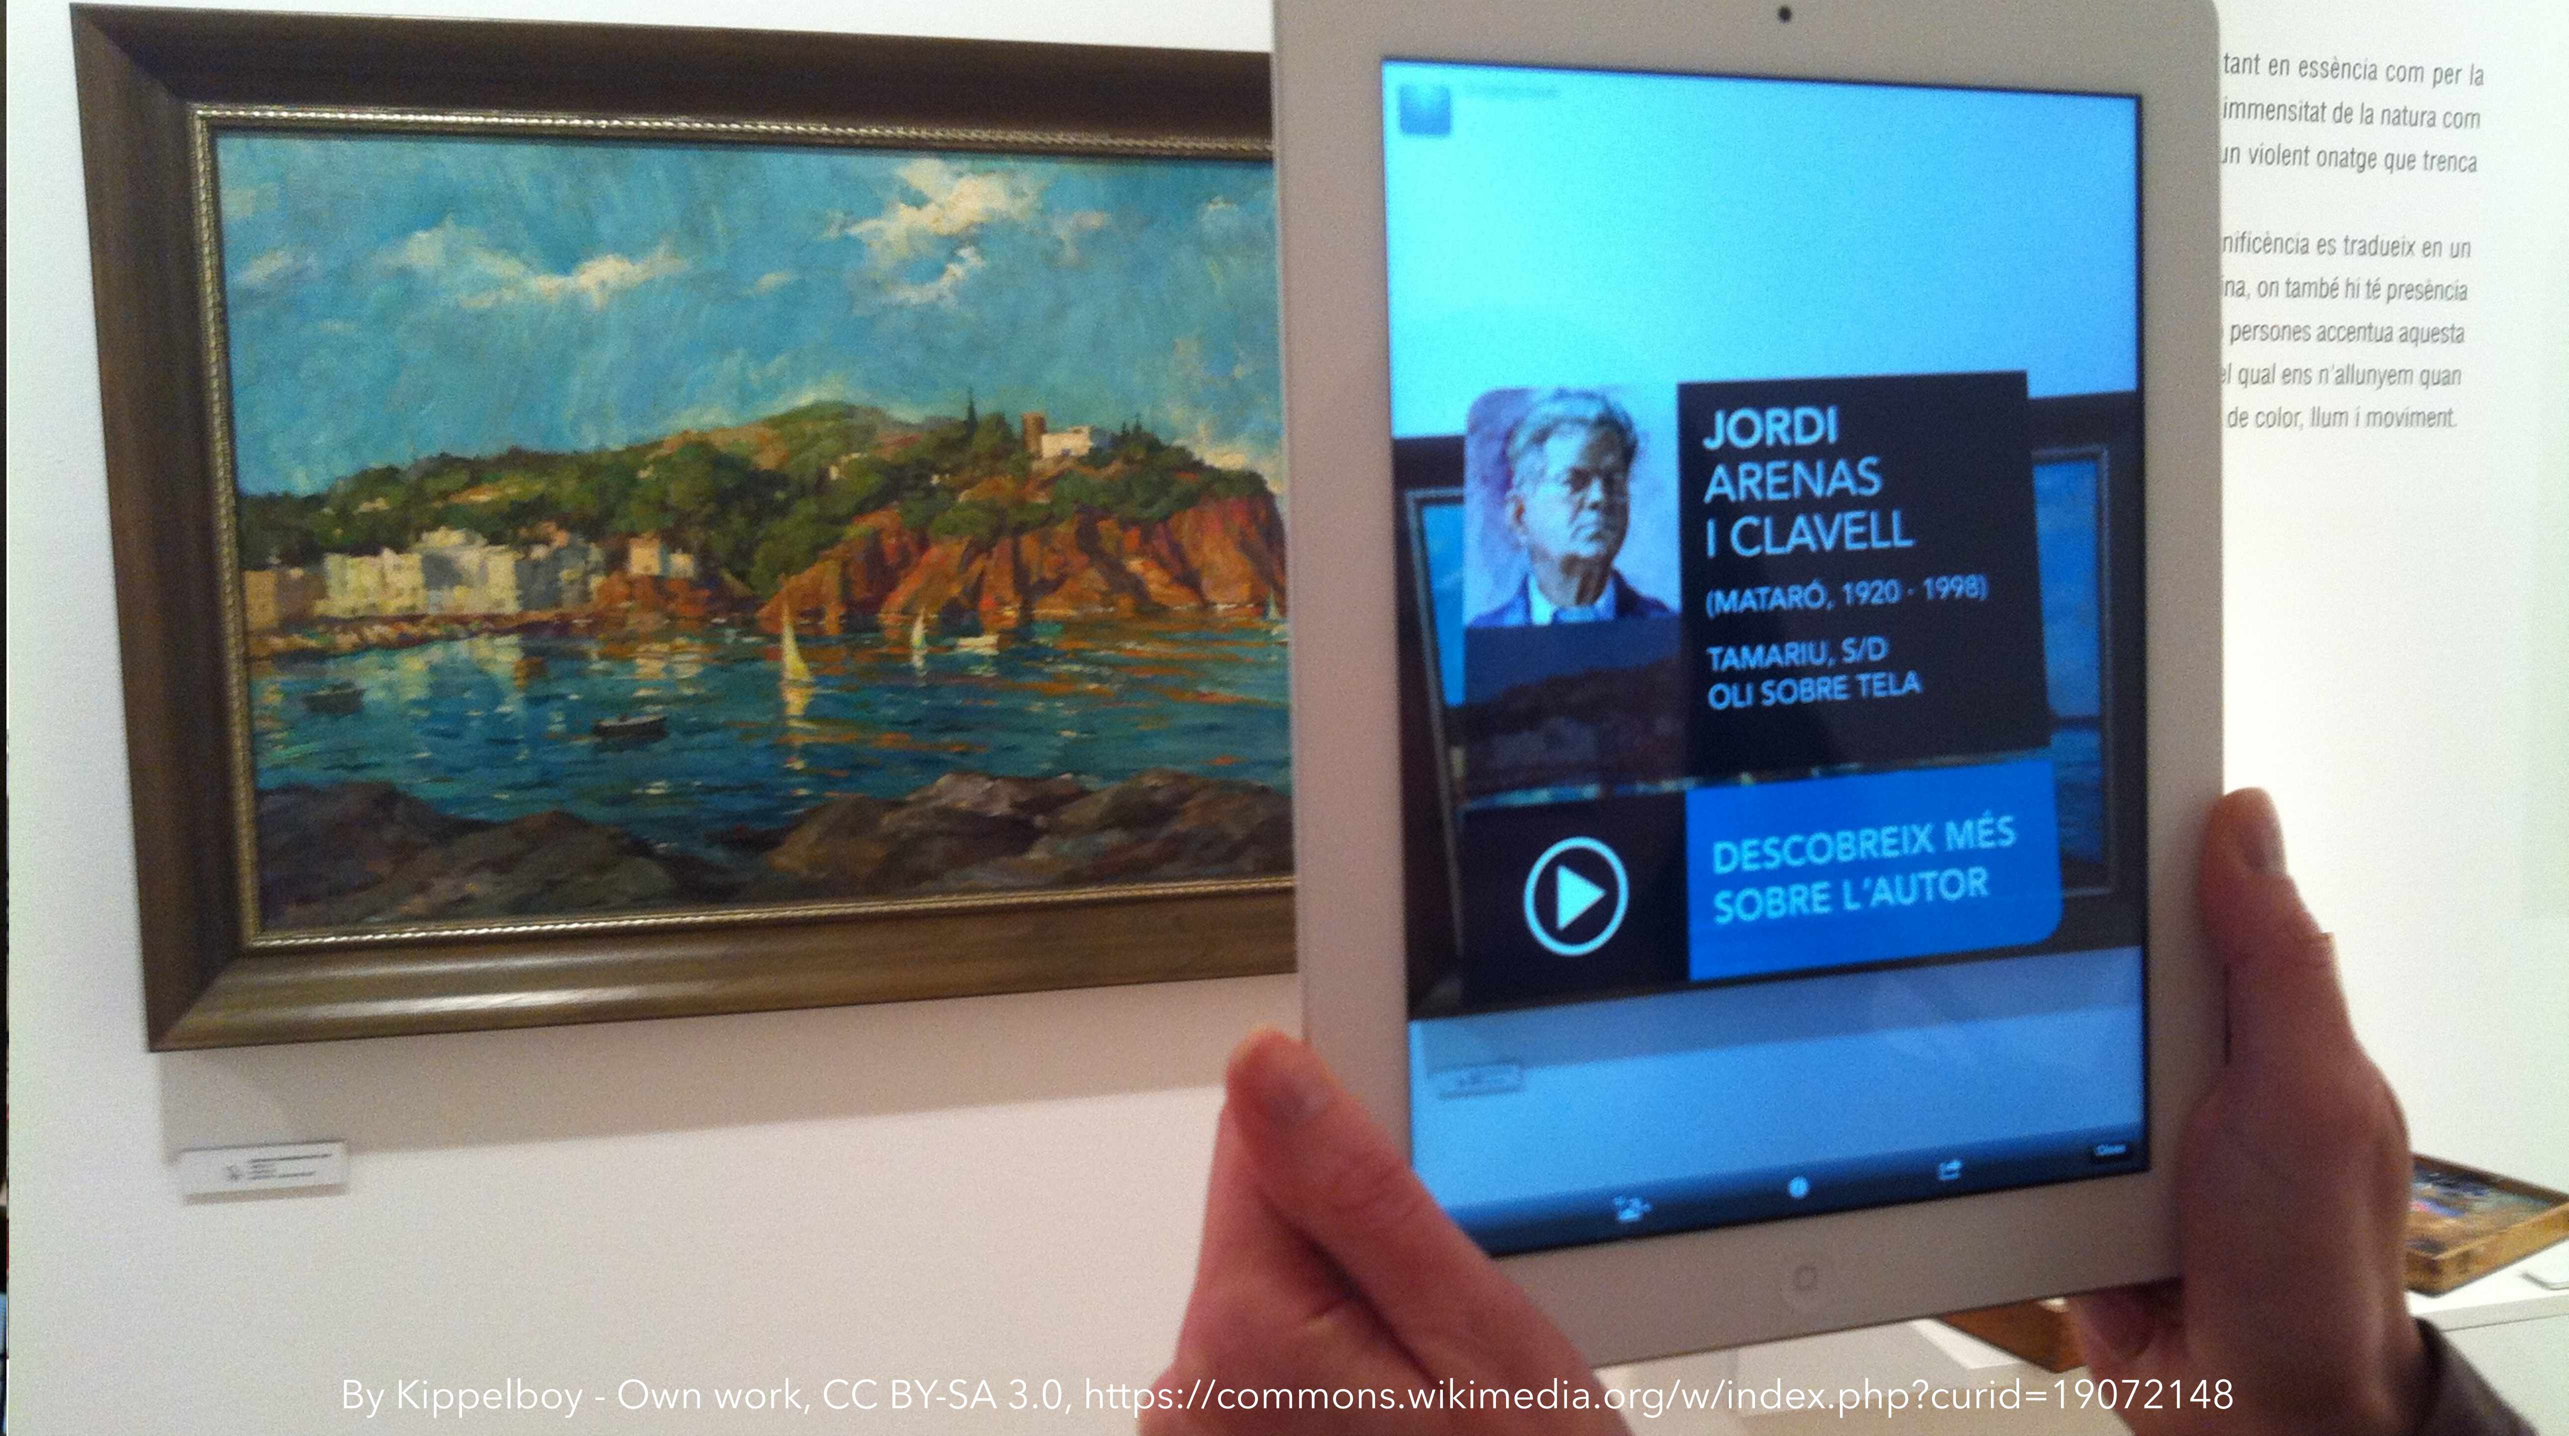
JORDI
ARENAS
I CLAVELL
(MATARÓ, 1920 - 1998)
TAMARIU, S/D
OLI SOBRE TELA
DESCOBREIX MÉS
SOBRE L'AUTOR


AA-B77 SP1
AA-B13 SP1
ocado


SCHOOL
SPEED
LIMIT
15
WAYMO

Background

What is "Computational Job Market Analysis"?

Our definition.

"The use of computational methods and tools to analyze trends, patterns, and dynamics within the job market."

What is “Computational Job Market Analysis”?

Our definition.

“The use of computational methods and tools to analyze trends, patterns, and dynamics within the job market.”

We leverage data science, machine learning, natural language processing, and other computational methods to gather, process, and interpret data related to job postings, employment statistics, labor market dynamics, skills demand, and other relevant factors.

What is "Computational Job Market Analysis"?

Our definition.

"The use of computational methods and tools to analyze trends, patterns, and dynamics within the job market."

We leverage data science, machine learning, natural language processing, and other computational methods to gather, process, and interpret data related to **job postings**, employment statistics, labor market dynamics, **skills** demand, and other relevant factors.

What are skills? (1)

Social scientist perspective.

- The concept of competence takes on diverse meanings across numerous fields and research domains (Rauner and Maclean, 2008, Zhao and Rauner, 2014):
  - A competence is having the necessary ability, authority, skill and knowledge. It is an individual's psychological character;
  - For competence to be defined, many factors play an important role, e.g., education systems, cultures, economical structures, vocational traditions as well as situations in the labour market.

What are skills? (2)

A taxonomical perspective.

- The European Skills/Competences, Qualifications & Occupations taxonomy (le Vrang et al., 2014) –> first public release in 2020
  - **13,890 concepts (knowledge, language skills, skills, transversal skills)**
  - 3,008 occupations
  - 28 languages

What are skills? (3)

A taxonomical perspective.

13,890 concepts (knowledge, language skills, skills, transversal skills)

- Two important “pillars”:
  - **Knowledge**: “knowledge means the outcome of the assimilation of information through learning” ≈ **Hard skills**
  - **Skill**: “the ability to apply knowledge and use know-how to complete tasks or solve problems” ≈ **Soft skills**
    - + **soft**, **transversal**, and **language** skills

Example

ESCO Illustrative Example.

European Commission

The European Skills, Competences, Qualifications and Occupations Taxonomy (ESCO)	
ESCO level 1: major group	S4 - Management skills

Example

ESCO Illustrative Example.

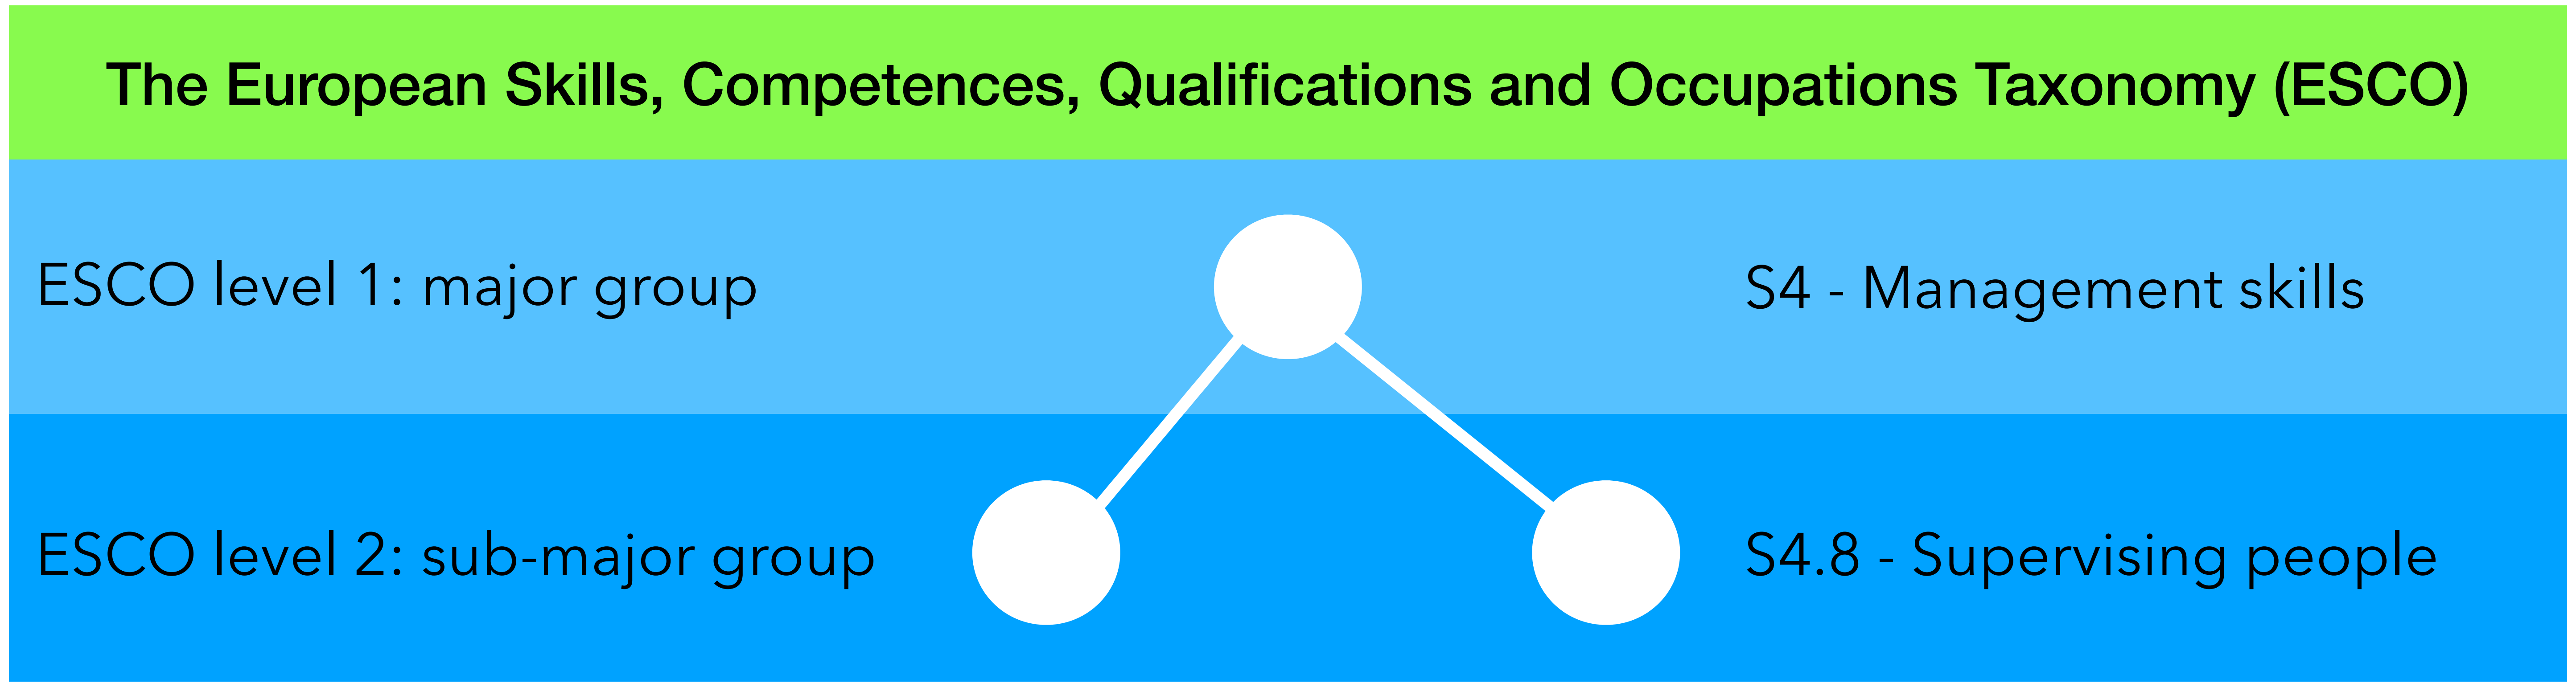

Example

ESCO Illustrative Example.

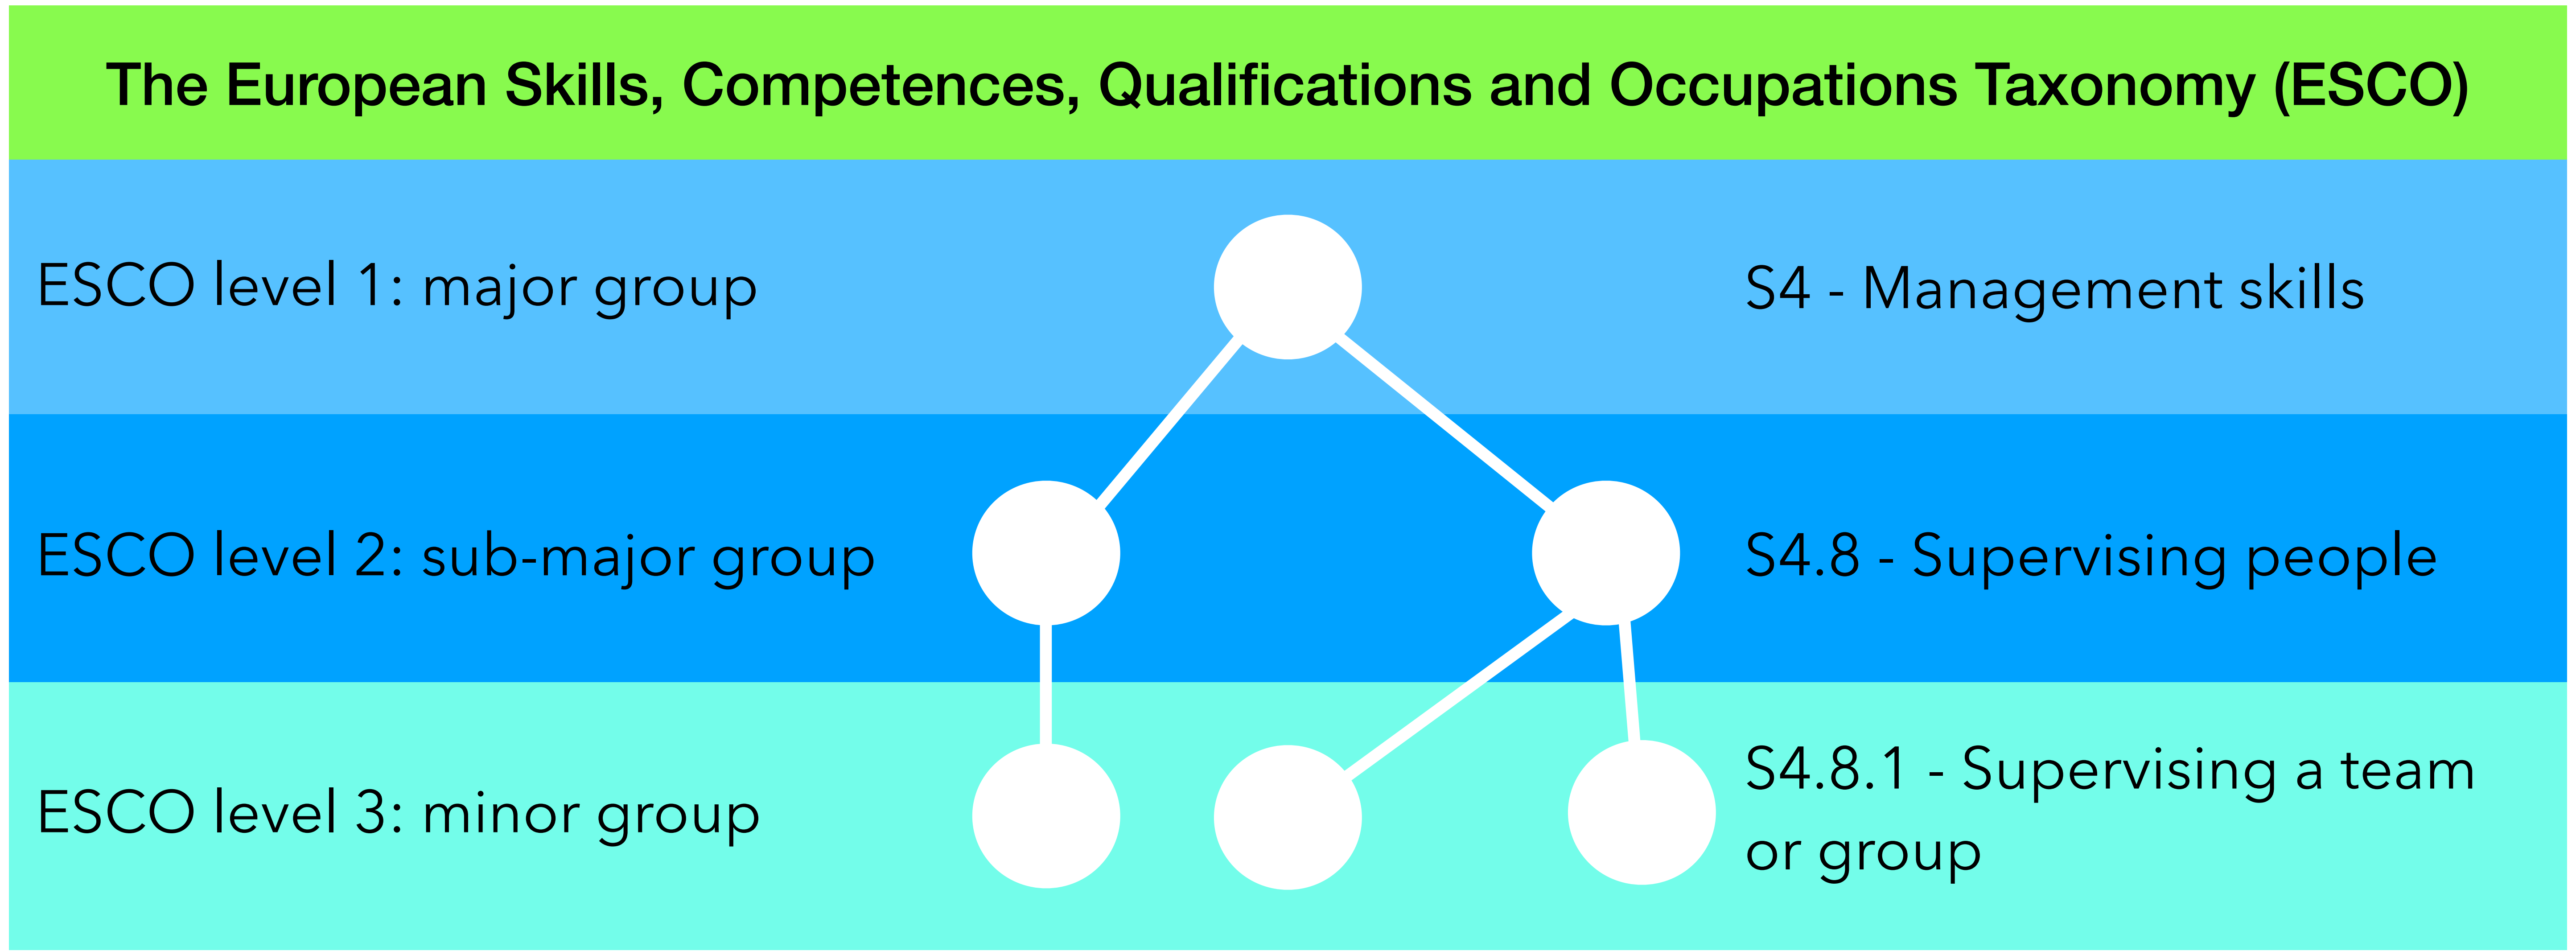

Example

ESCO Illustrative Example.

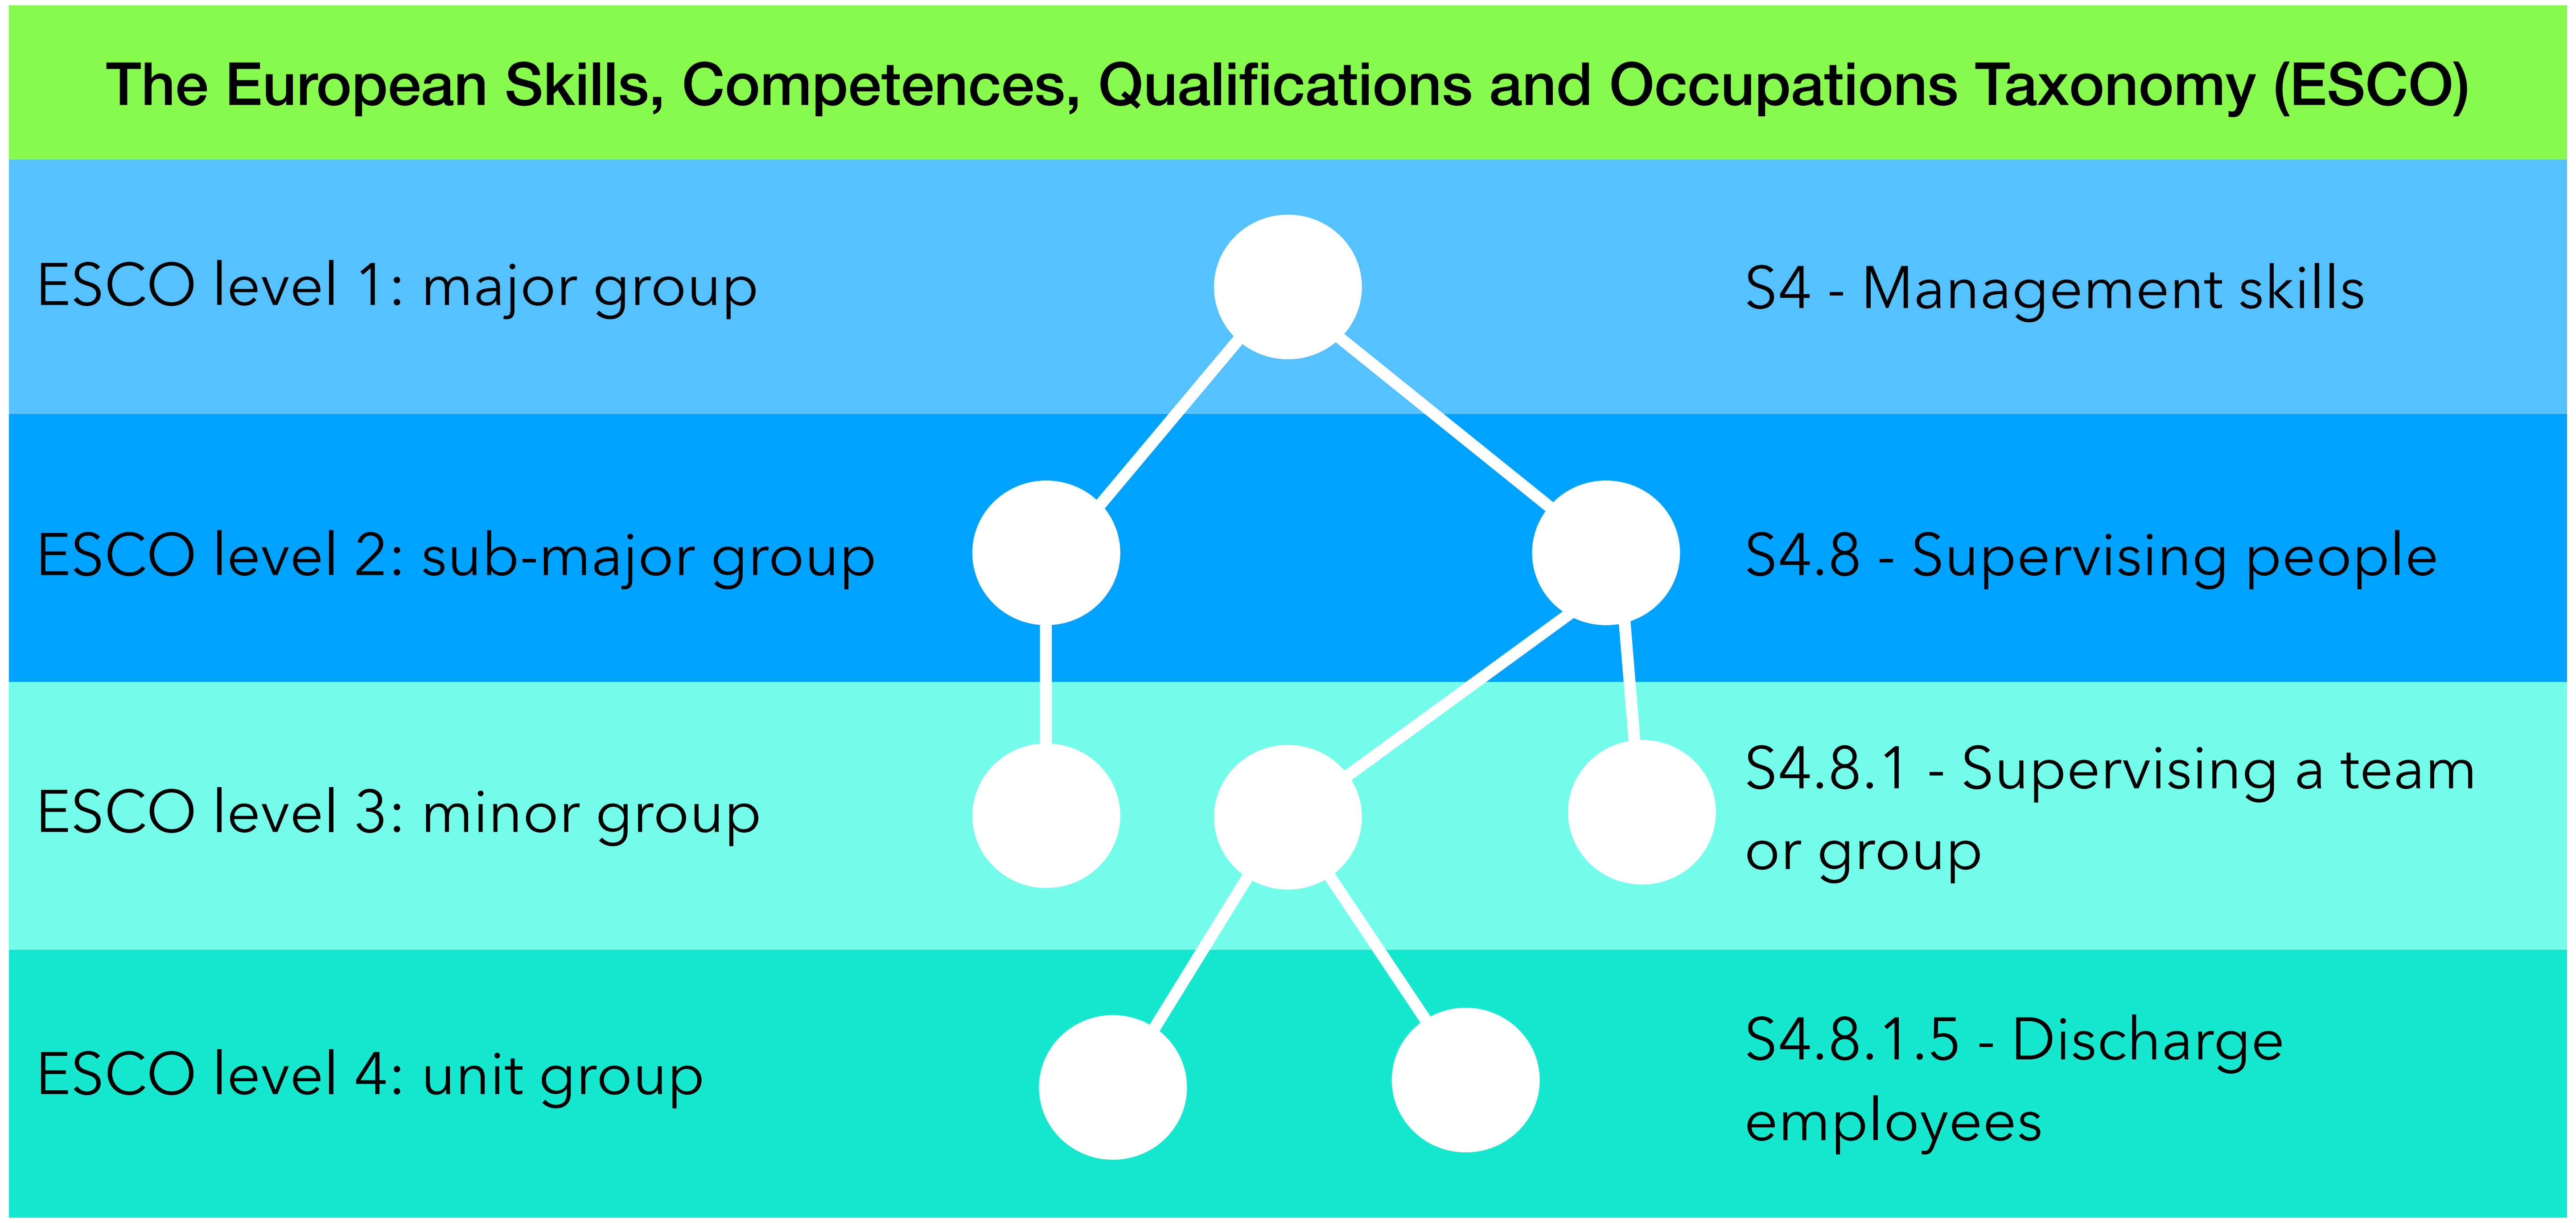

Example

ESCO Illustrative Example.

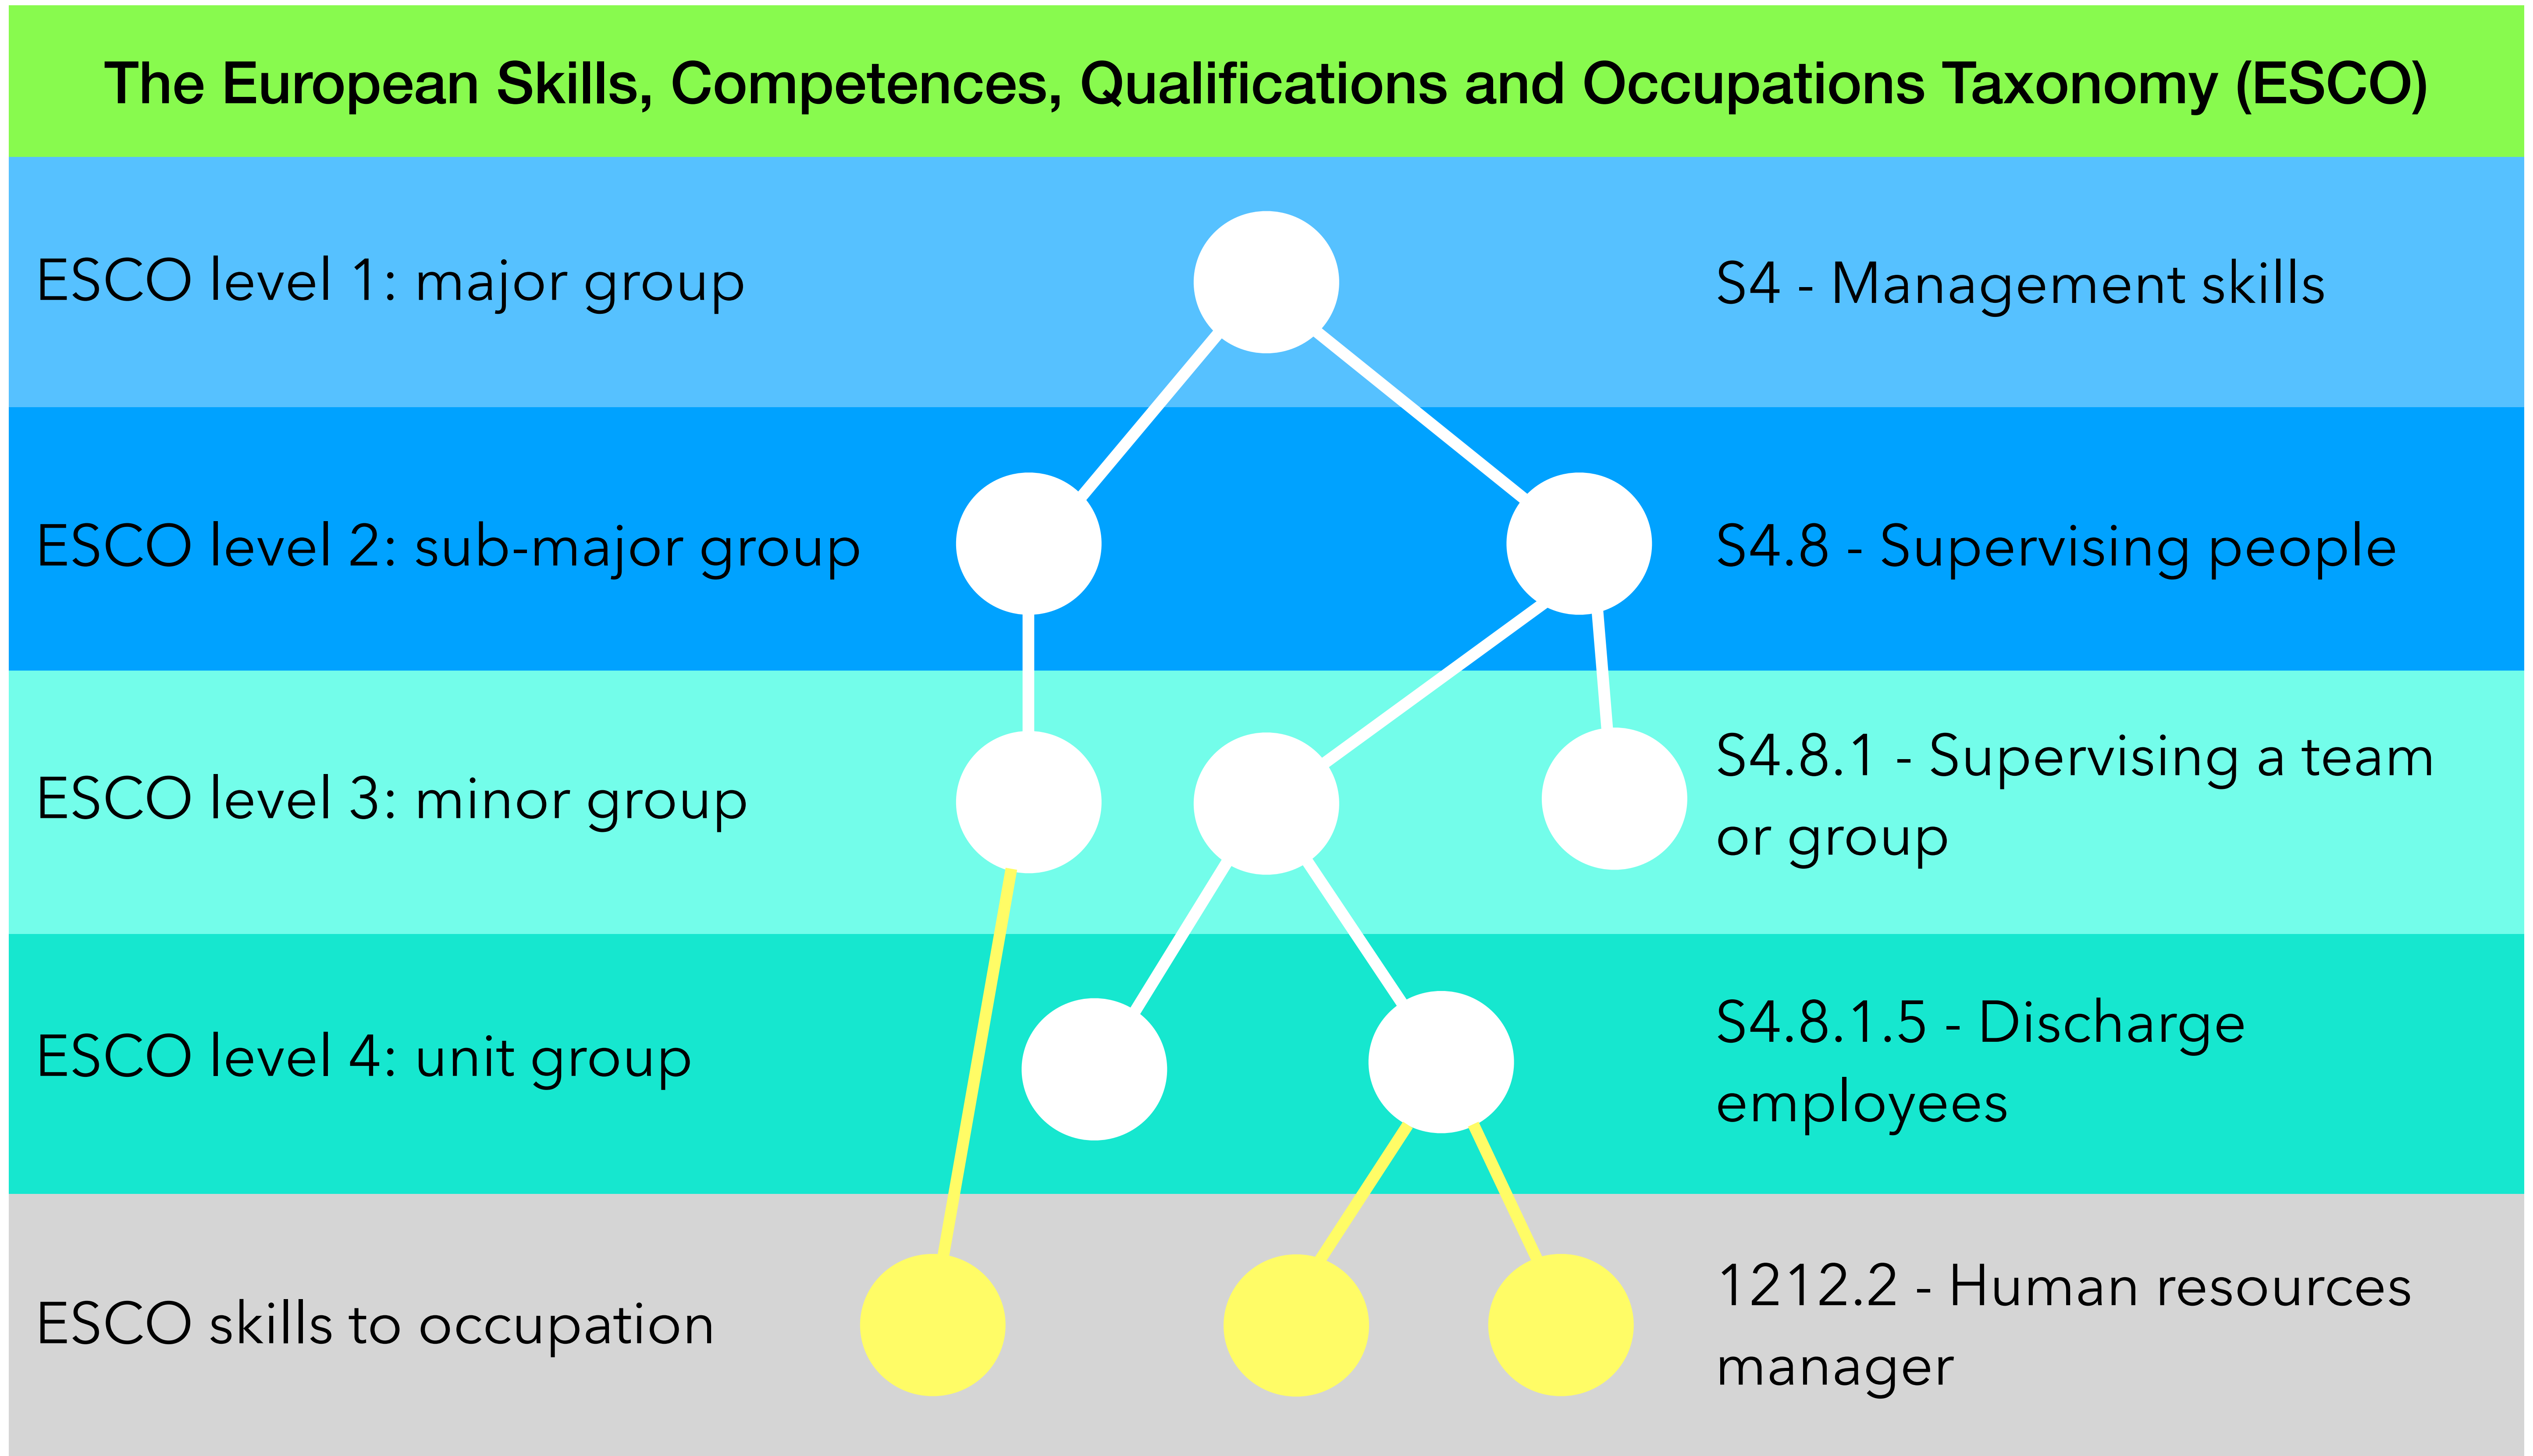

Setting the NLP context (1)

Skill Extraction == Sequence Labeling.

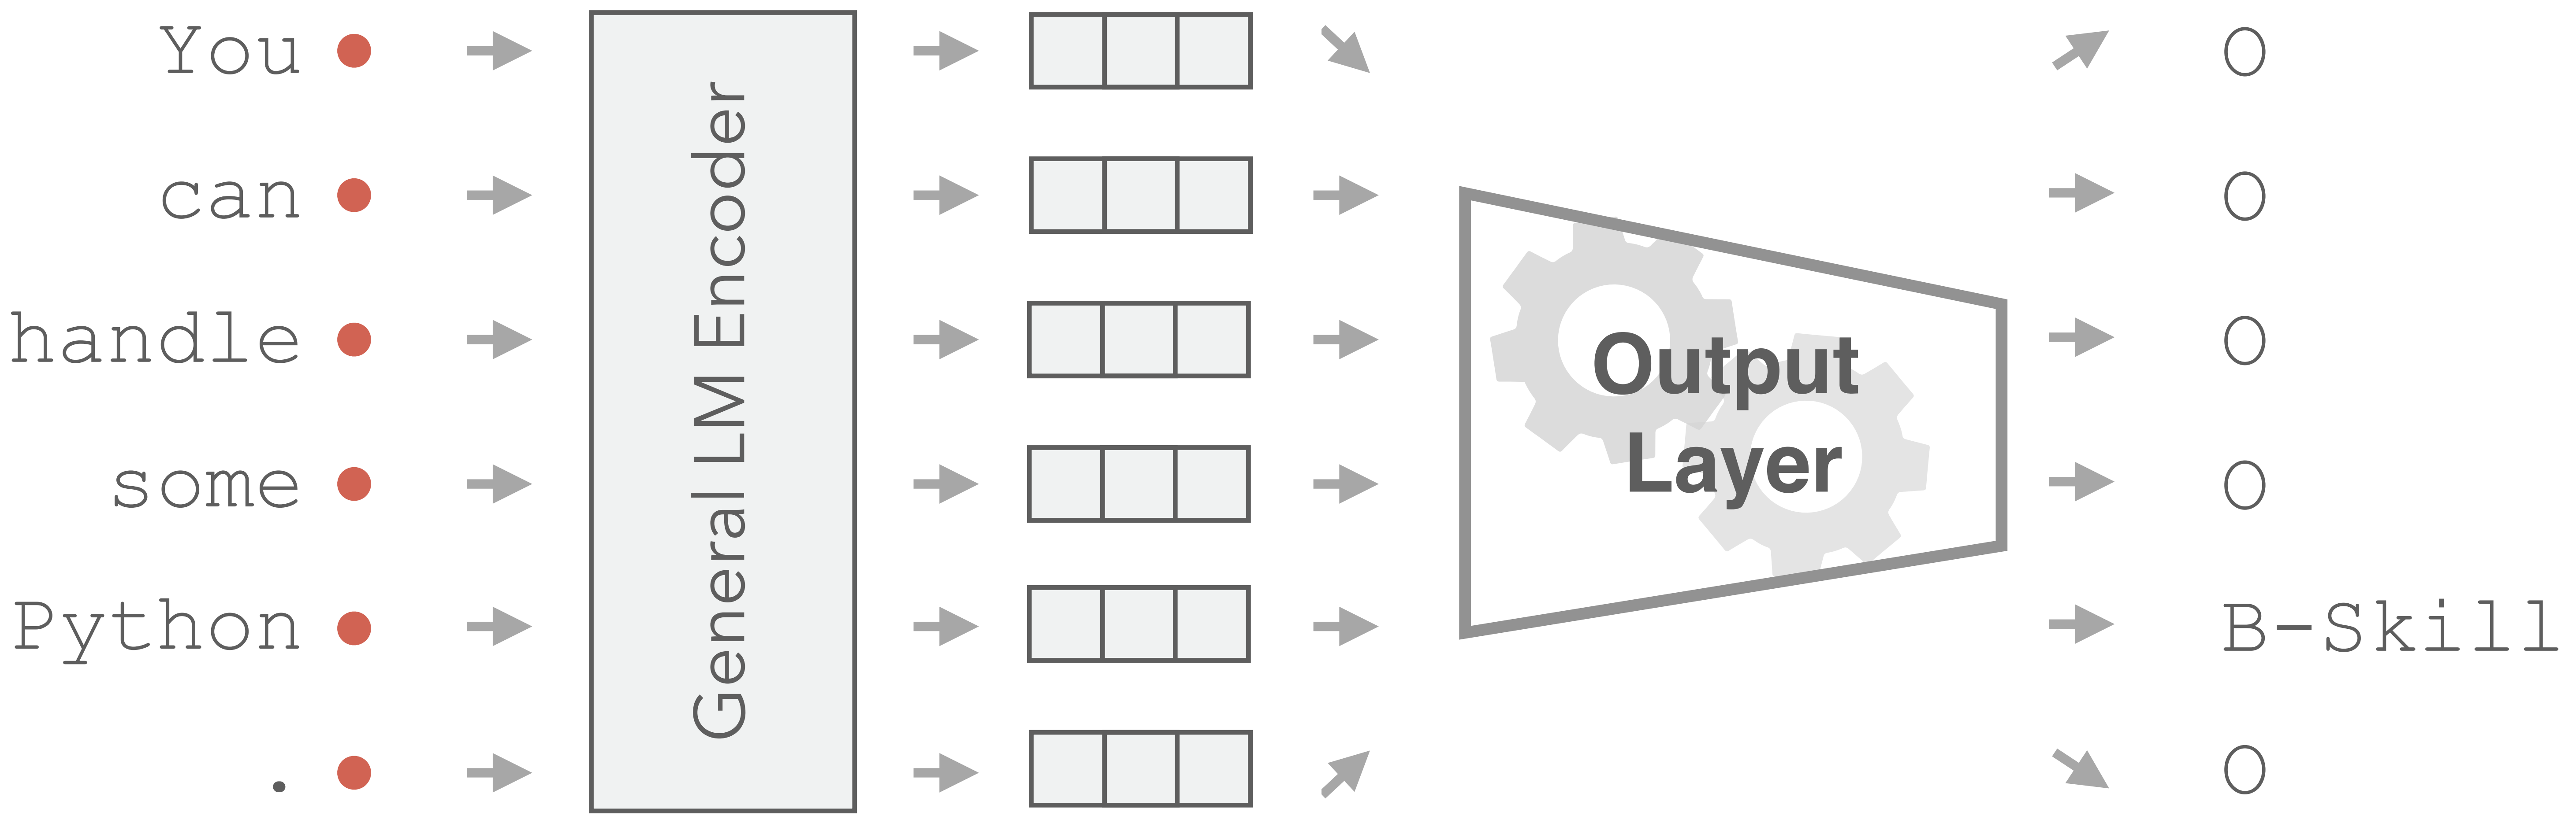

Setting the NLP context (1)

Skill Extraction == Sequence Labeling.

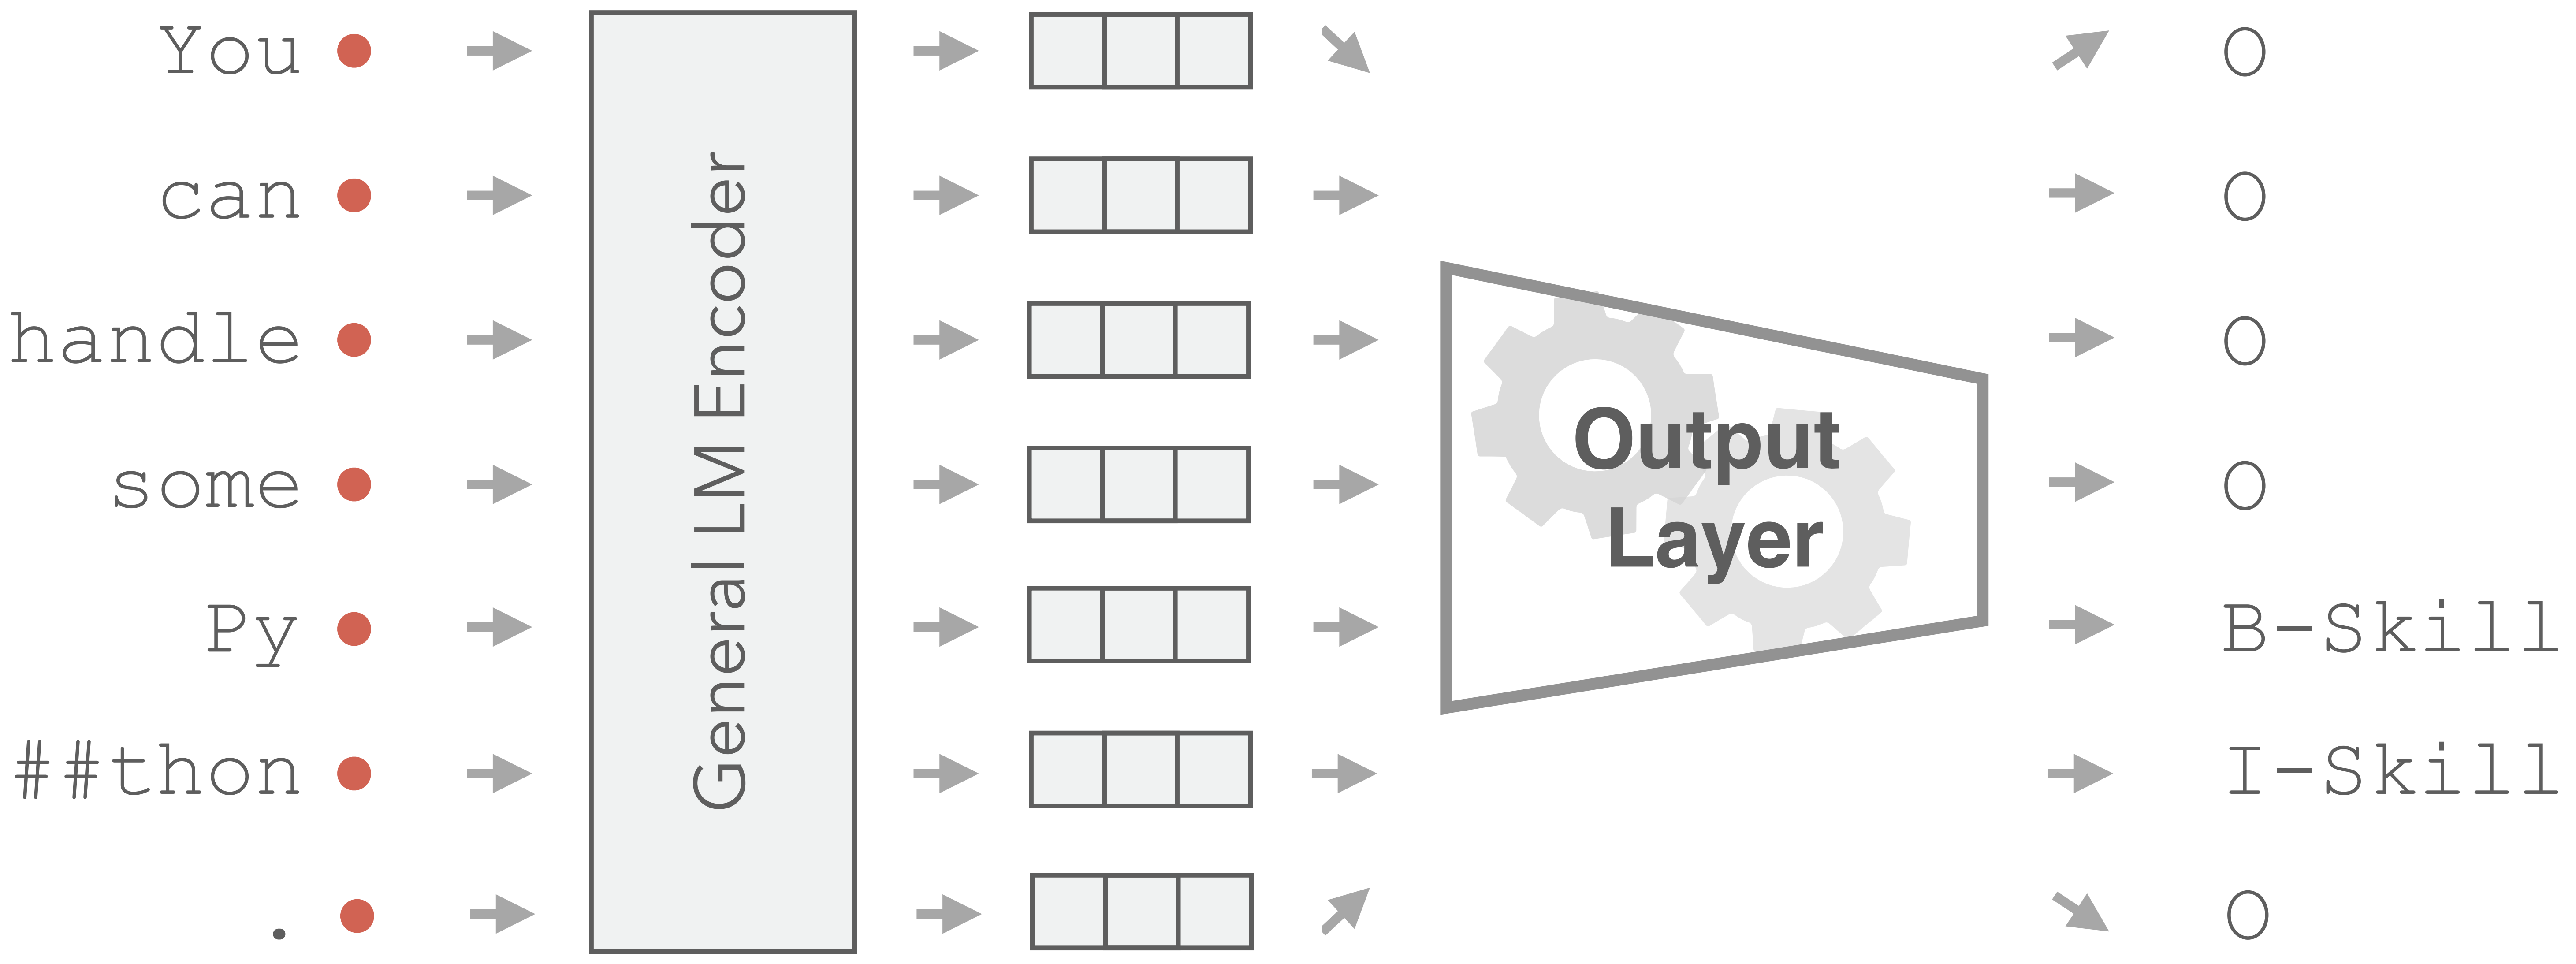

Setting the NLP context (2)

Skill Extraction == Sequence Labeling

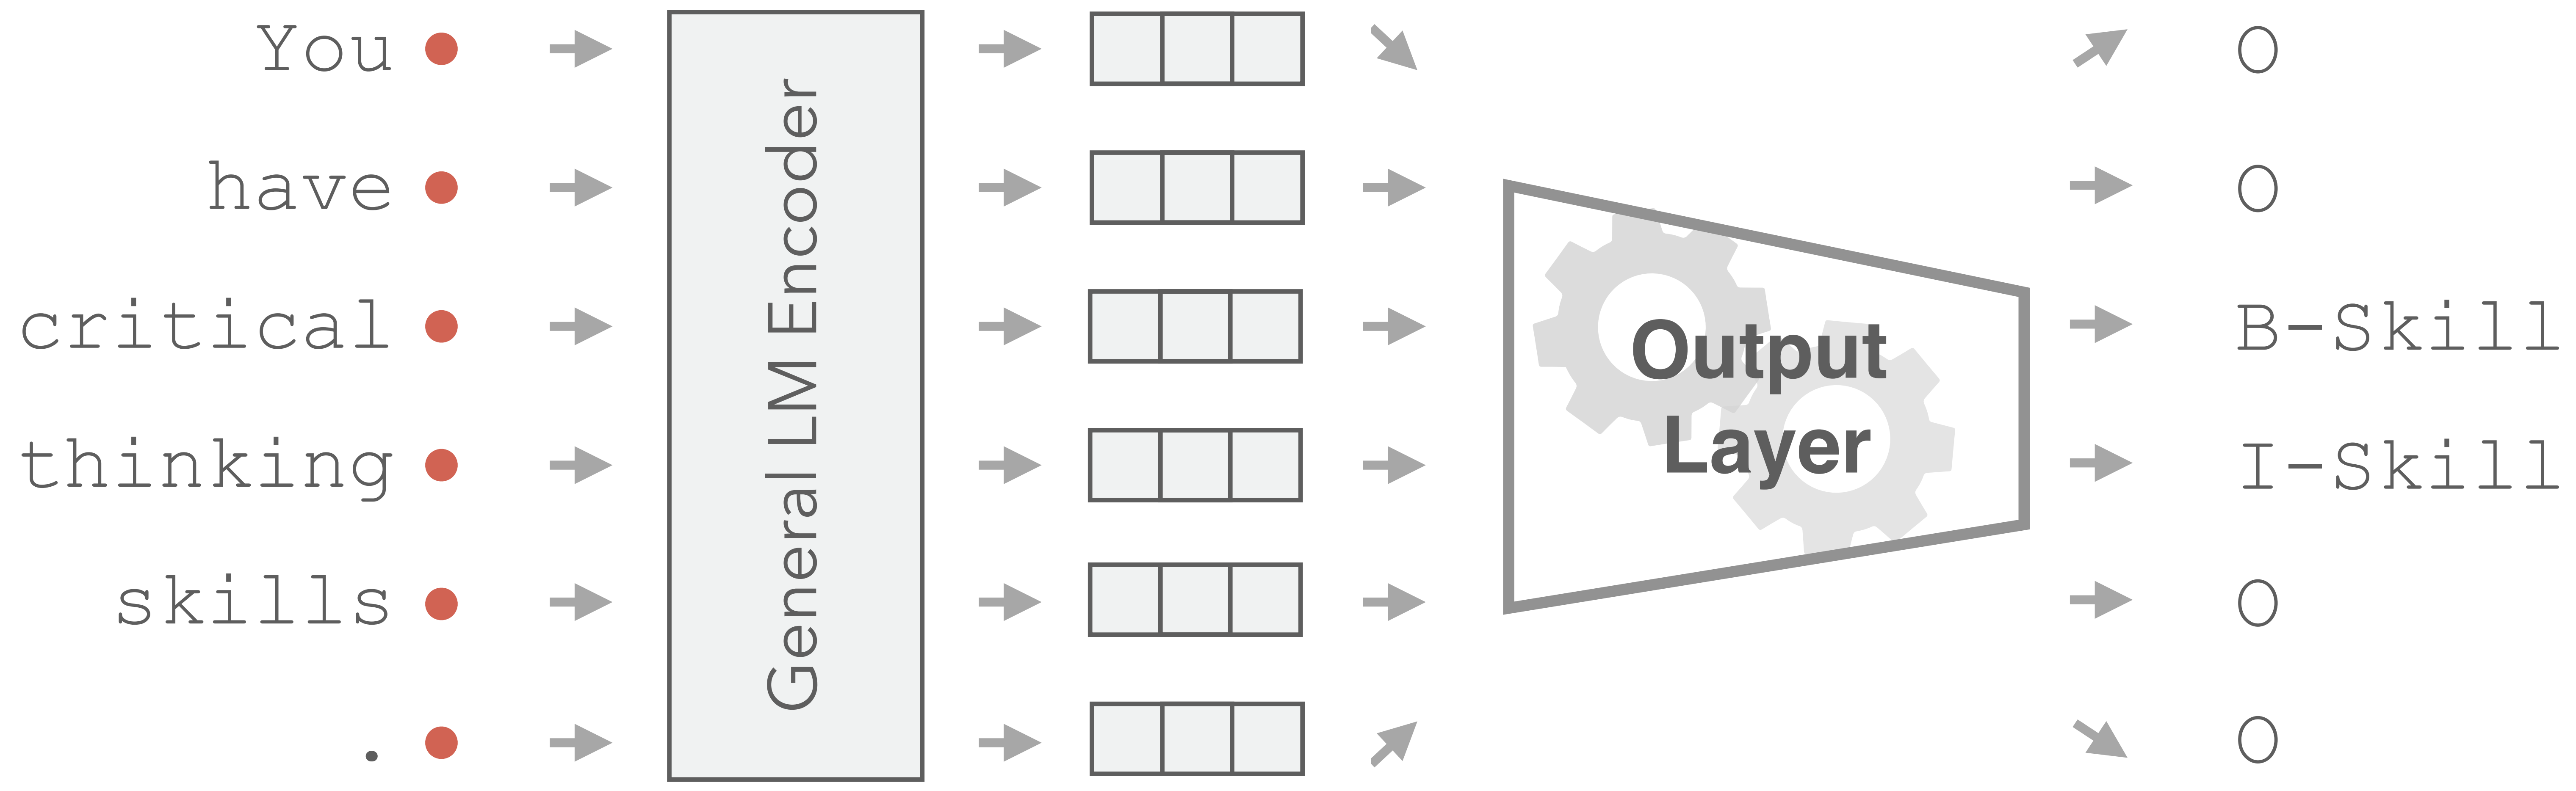

Challenges

Challenges

Five Challenges with Skills.

1. **Implicitness**: The skill "teamwork" may manifest in various formats such as "being able to work together," "working in a team," or "collaborating efficiently."
2. **Span Length**: Skills, by their nature, can exhibit considerable length, posing a challenge in extraction with NLP models.
3. **Long-tail**: Certain skills are infrequently mentioned in a comprehensive set of job postings.
4. **Discontinuity**: There is a discontinuity of skill spans, analogous to what is referred to as *coordinated conjunctions* in linguistics.
5. **Output Space**: If we use a taxonomy like ESCO as the labels, this results in a substantial output space during inference, adding a layer of complexity to the modeling process.

Challenges

Five Challenges with Skills.

1. **Implicitness**: The skill “teamwork” may manifest in various formats such as “being able to work together,” “working in a team,” or “collaborating efficiently.”
2. **Span Length**: Skills, by their nature, can exhibit considerable length, posing a challenge in extraction with NLP models.
3. **Long-tail**: Certain skills are infrequently mentioned in a comprehensive set of job postings.
4. **Discontinuity**: There is a discontinuity of skill spans, analogous to what is referred to as *coordinated conjunctions* in linguistics.
5. **Output Space**: If we use a taxonomy like ESCO as the labels, this results in a substantial output space during inference, adding a layer of complexity to the modeling process.

Computational Job Market Analysis Storyline

The red thread through the thesis.

Data

JobStack
(Jensen et al., 2021, NoDaLiDa)

CAL
(Zhang and Plank, 2021, EMNLP)

SkillSpan
(Zhang et al., 2022, NAACL)

Kompetencer
(Zhang et al., 2022, LREC)

→

Modeling

WeakSupervision
(Zhang et al., 2022, RecSysHR)

ESCOXLM-R
(Zhang et al., 2023, ACL)

NNOSE
(Zhang et al., 2024, EACL)

→

Linking

EL Job Market
(Zhang et al., 2024, EACL)

Computational Job Market Analysis Storyline

The red thread through the thesis.

Data		Modeling		Linking
JobStack (Jensen et al., 2021, NoDaLiDa)	→	WeakSupervision (Zhang et al., 2022, RecSysHR)	→	**EL Job Market (Zhang et al., 2024, EACL)**
CAL (Zhang and Plank, 2021, EMNLP)		ESCOXLM-R (Zhang et al., 2023, ACL)		
SkillSpan (Zhang et al., 2022, NAACL)		**NNOSE (Zhang et al., 2024, EACL)**		
Kompetencer (Zhang et al., 2022, LREC)				

Getting Annotated Data

Getting Annotated Data

Why?

- We need annotated job posting data to train supervised models for skill extraction.
- Approaches:
  - Existing datasets (2 during this time);
  - Scraping + Annotating yourself (time-consuming).

Getting Annotated Data

Why?

- We need annotated job posting data to train supervised models for skill extraction.
- Approaches:
  - Existing datasets (2 during this time);
  - **Scraping + Annotating yourself (time-consuming).**

SkillSpan

  

SkillSpan: Hard and Soft Skill Extraction from English Job Postings.

Mike **Zhang**, Kristian Nørgaard **Jensen**, Sif Dam **Sonniks**, and Barbara **Plank.**

In Proceedings of NAACL 2022.

Sources

Where does it come from?

- We have job postings from three sources:
  - **Tech**: Stackoverflow job postings;
  - **Big**: Indeed GB (not public);
  - **House**: Jobnet (thanks to STAR!)

Quick Recap ESCO

Knowledge and Skills

- **13,890 concepts (knowledge, language skills, skills, transversal skills)**
- Two important “pillars”:
  - **Knowledge**: “knowledge means the outcome of the assimilation of information through learning” **≈ Hard skills**
  - **Skill**: “the ability to apply knowledge and use know-how to complete tasks or solve problems” **≈ Soft skills**
    - + **soft**, **transversal**, and **language** skills

Example (1): Skills

SkillSpan: Hard and Soft Skill Extraction from English Job Postings.

"The successful candidate will be self-motivated and capable of working on their own initiative, an excellent communicator, with both customers and our highly motivated team at the company."

Knowledge

Skill

Example (2): Soft-like skills

SkillSpan: Hard and Soft Skill Extraction from English Job Postings.

> “We are looking for someone who is a team player with a reliable approach to problem solving.”

Knowledge

Skill

Example (3): Knowledge and nested components

SkillSpan: Hard and Soft Skill Extraction from English Job Postings.

"Experience in the kitchen industry and competent with computers. Strong experience in analyzing supply chain strategies."

Knowledge

Skill

Example (4): Long skills

Challenge 2: Long skills

"Ability to read, write, configure, design, and script end-to-end service telemetry alerting and self-healing capabilities for platforms."

Knowledge

Skill

Key Statistics

Dataset Statistics.

- 391 job postings:
  - 14.5K sentences;
    - 232K tokens;
  - 12.5K annotations;
    - 300 overlapping spans.
- 8 months of hard work.

Recap

Skill Extraction == Sequence Labeling

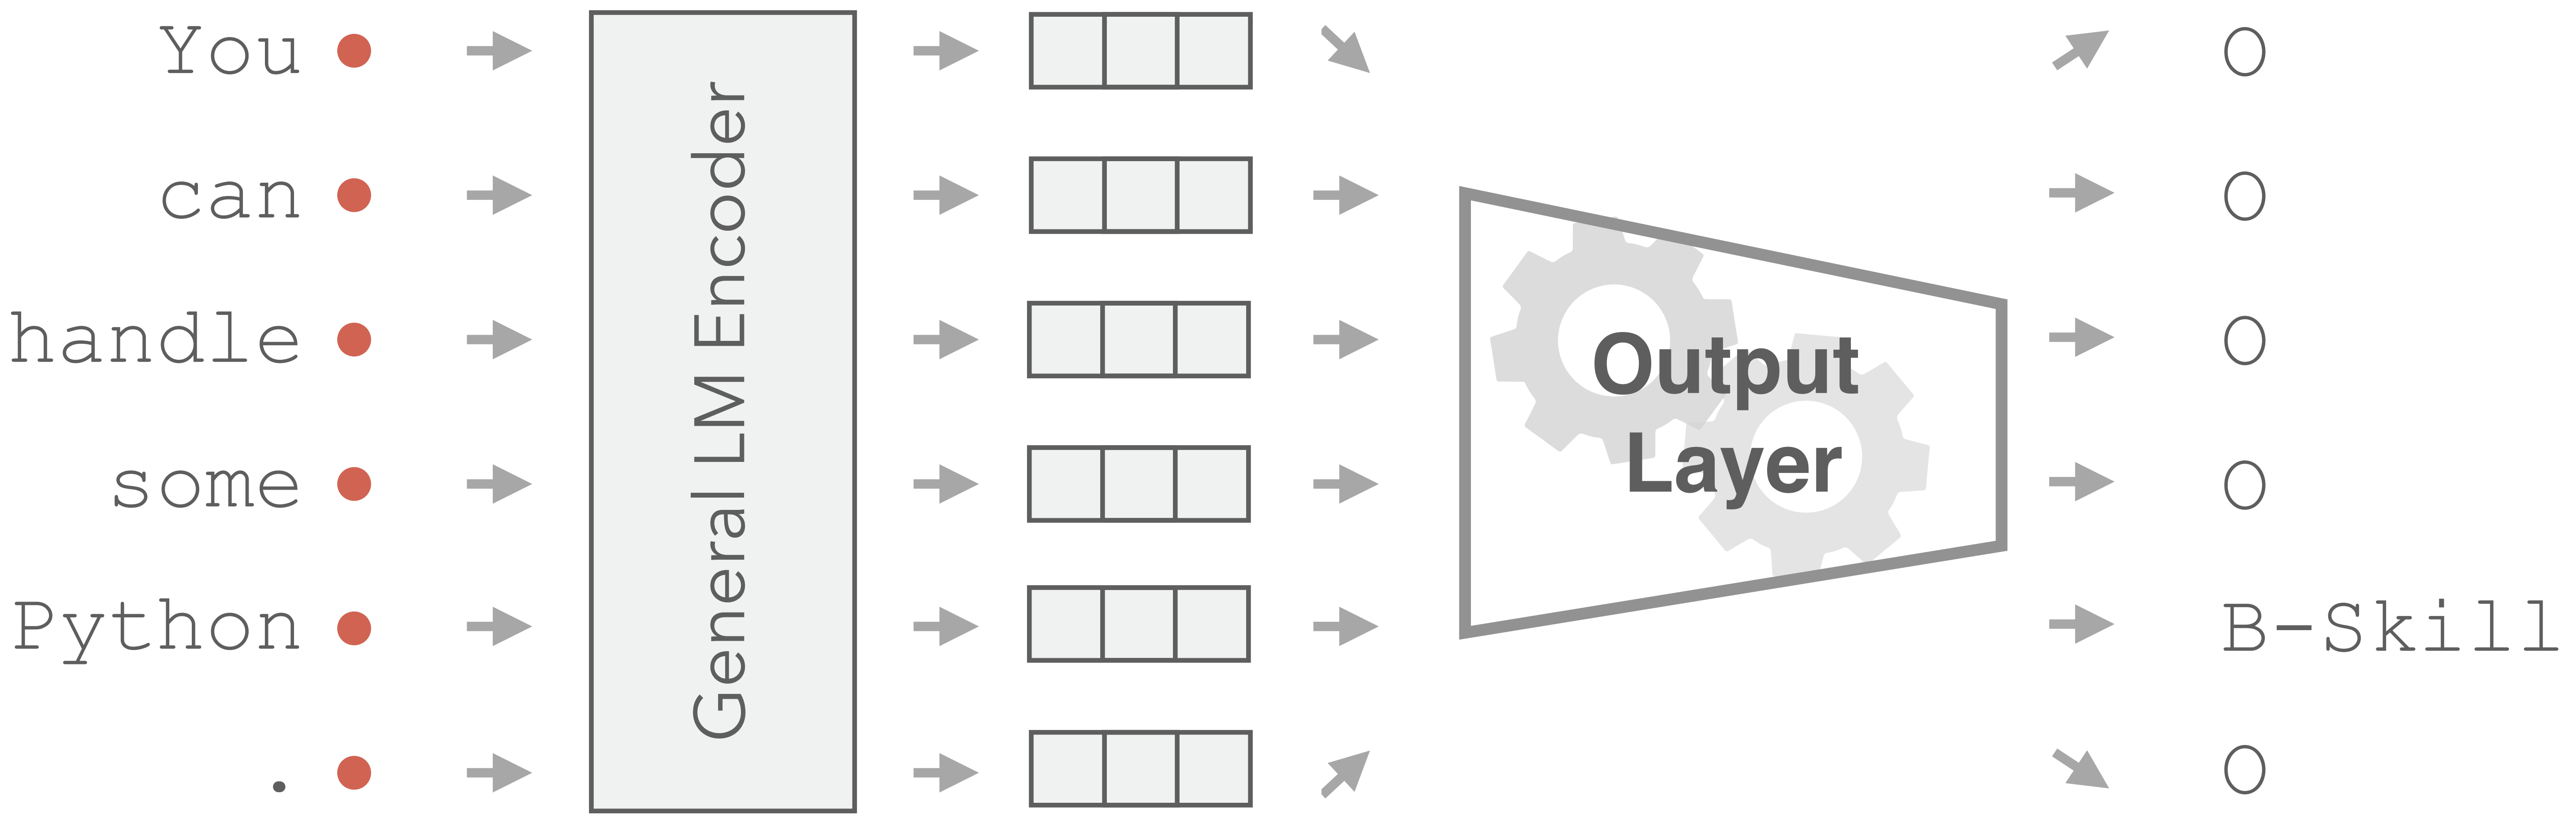

Two sets of experiments

SkillSpan

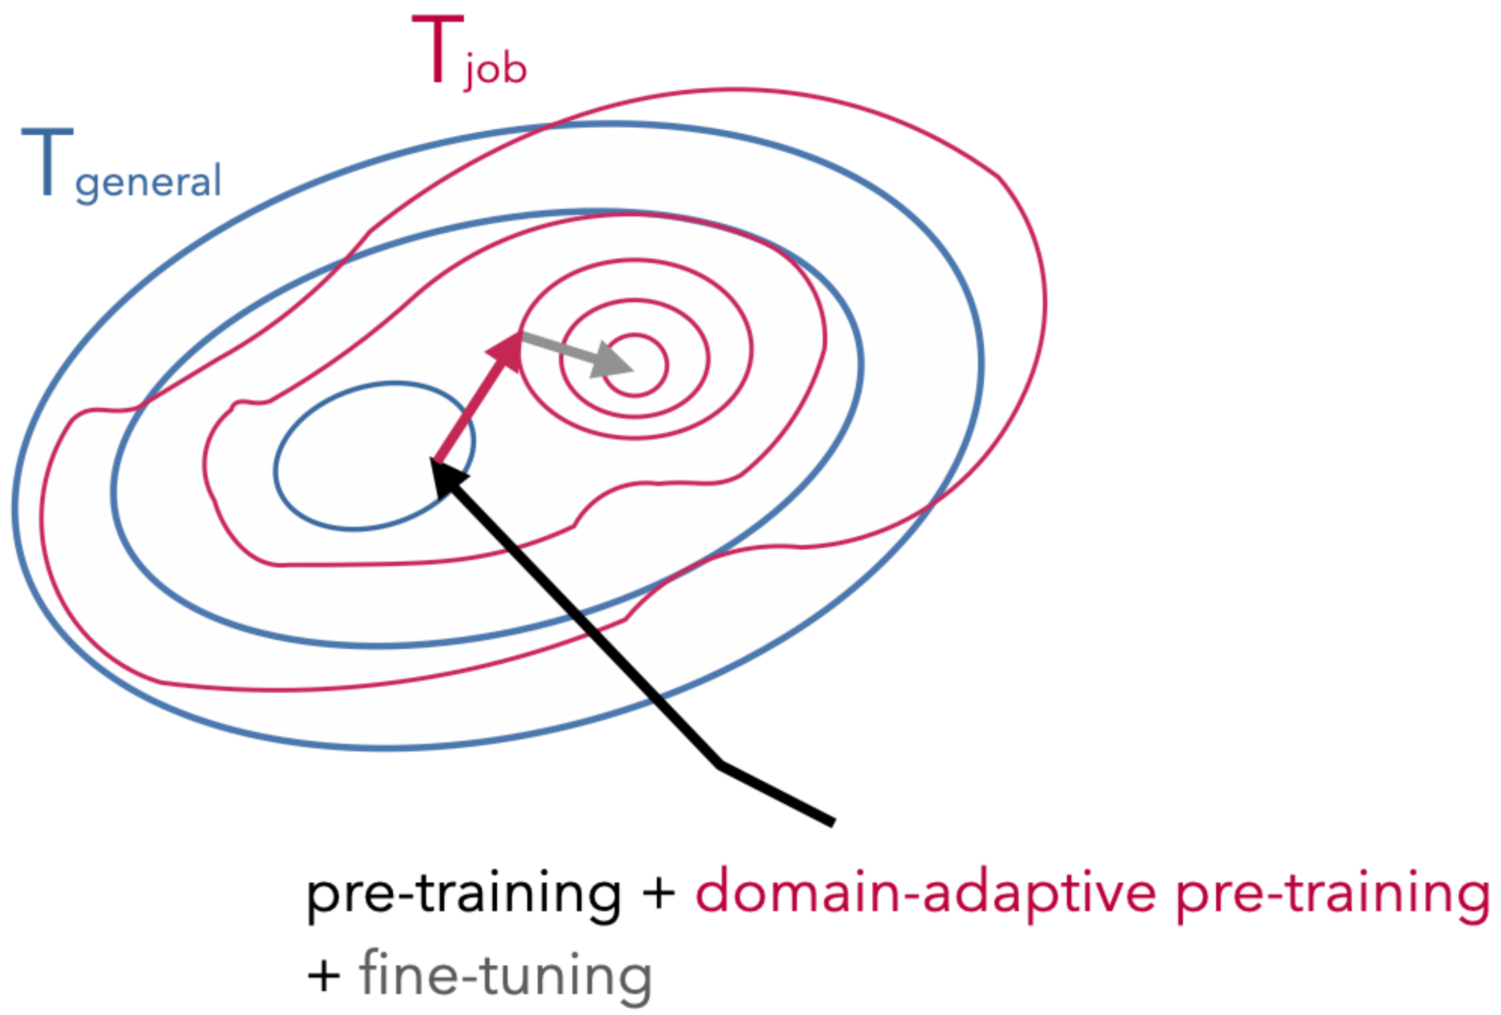

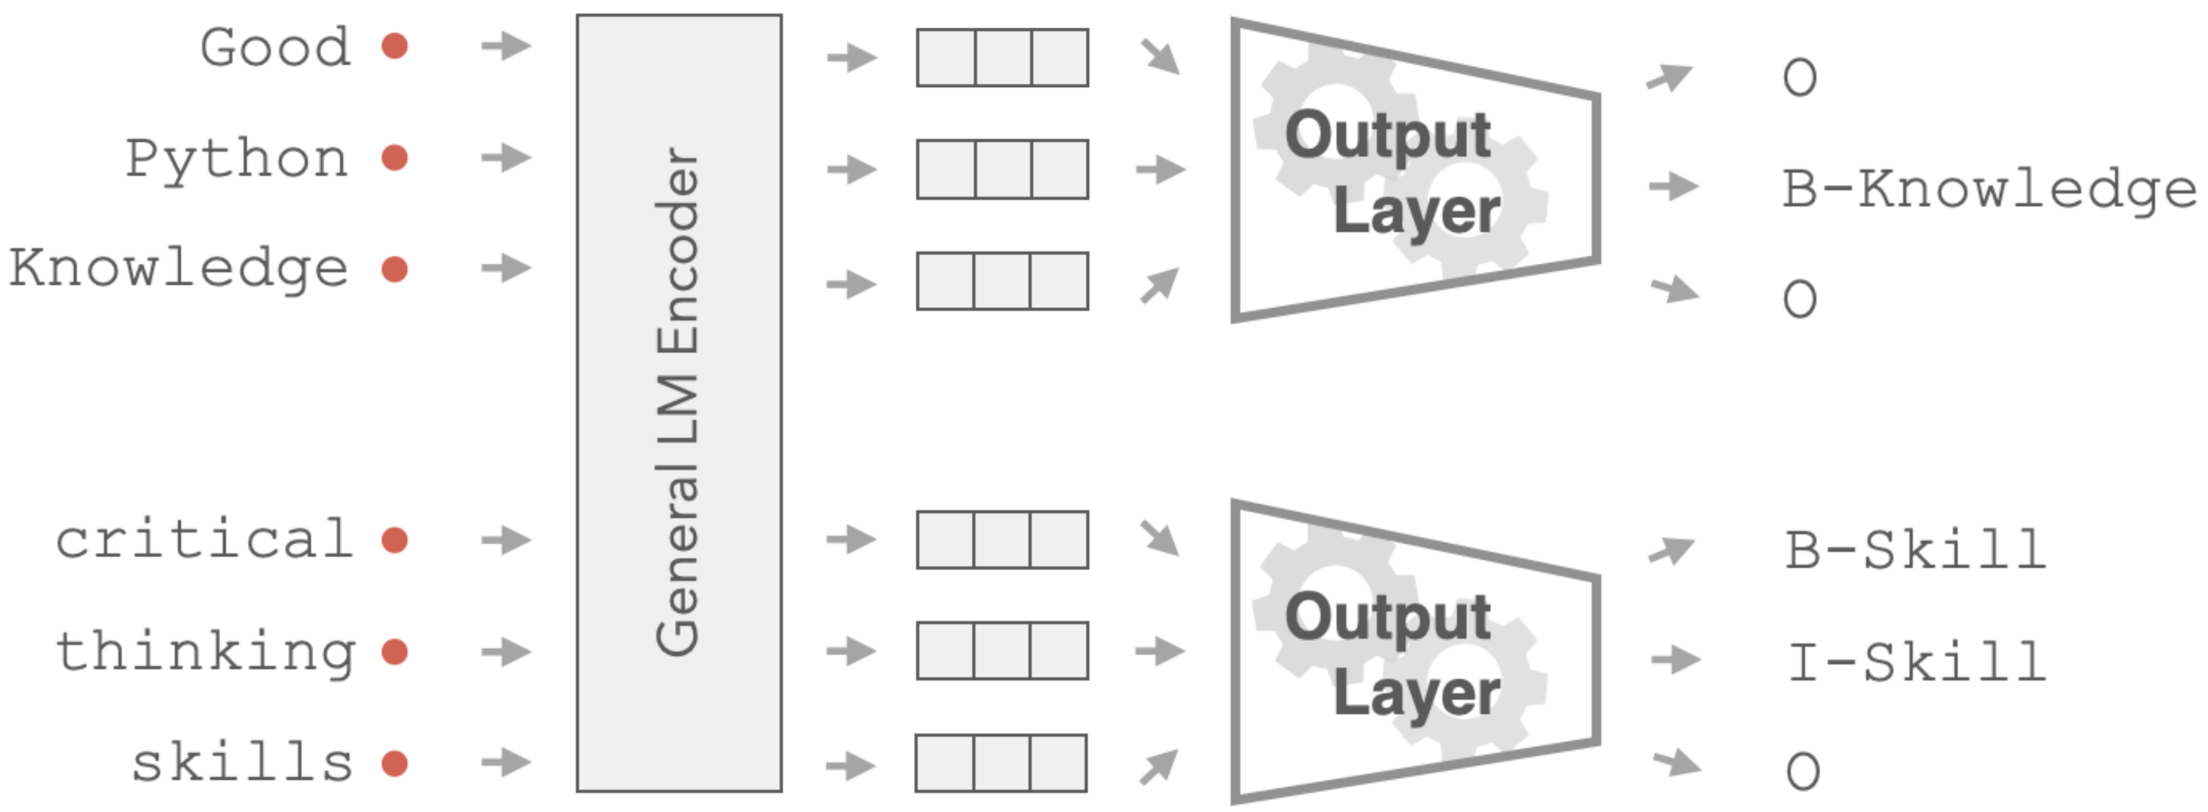

Domain-adaptive pretraining (Gururangan et al., 2020)

Multi-task Learning (Caruana, 1997)

Domain-adaptive Pre-training

Gururangan et al. (2020)

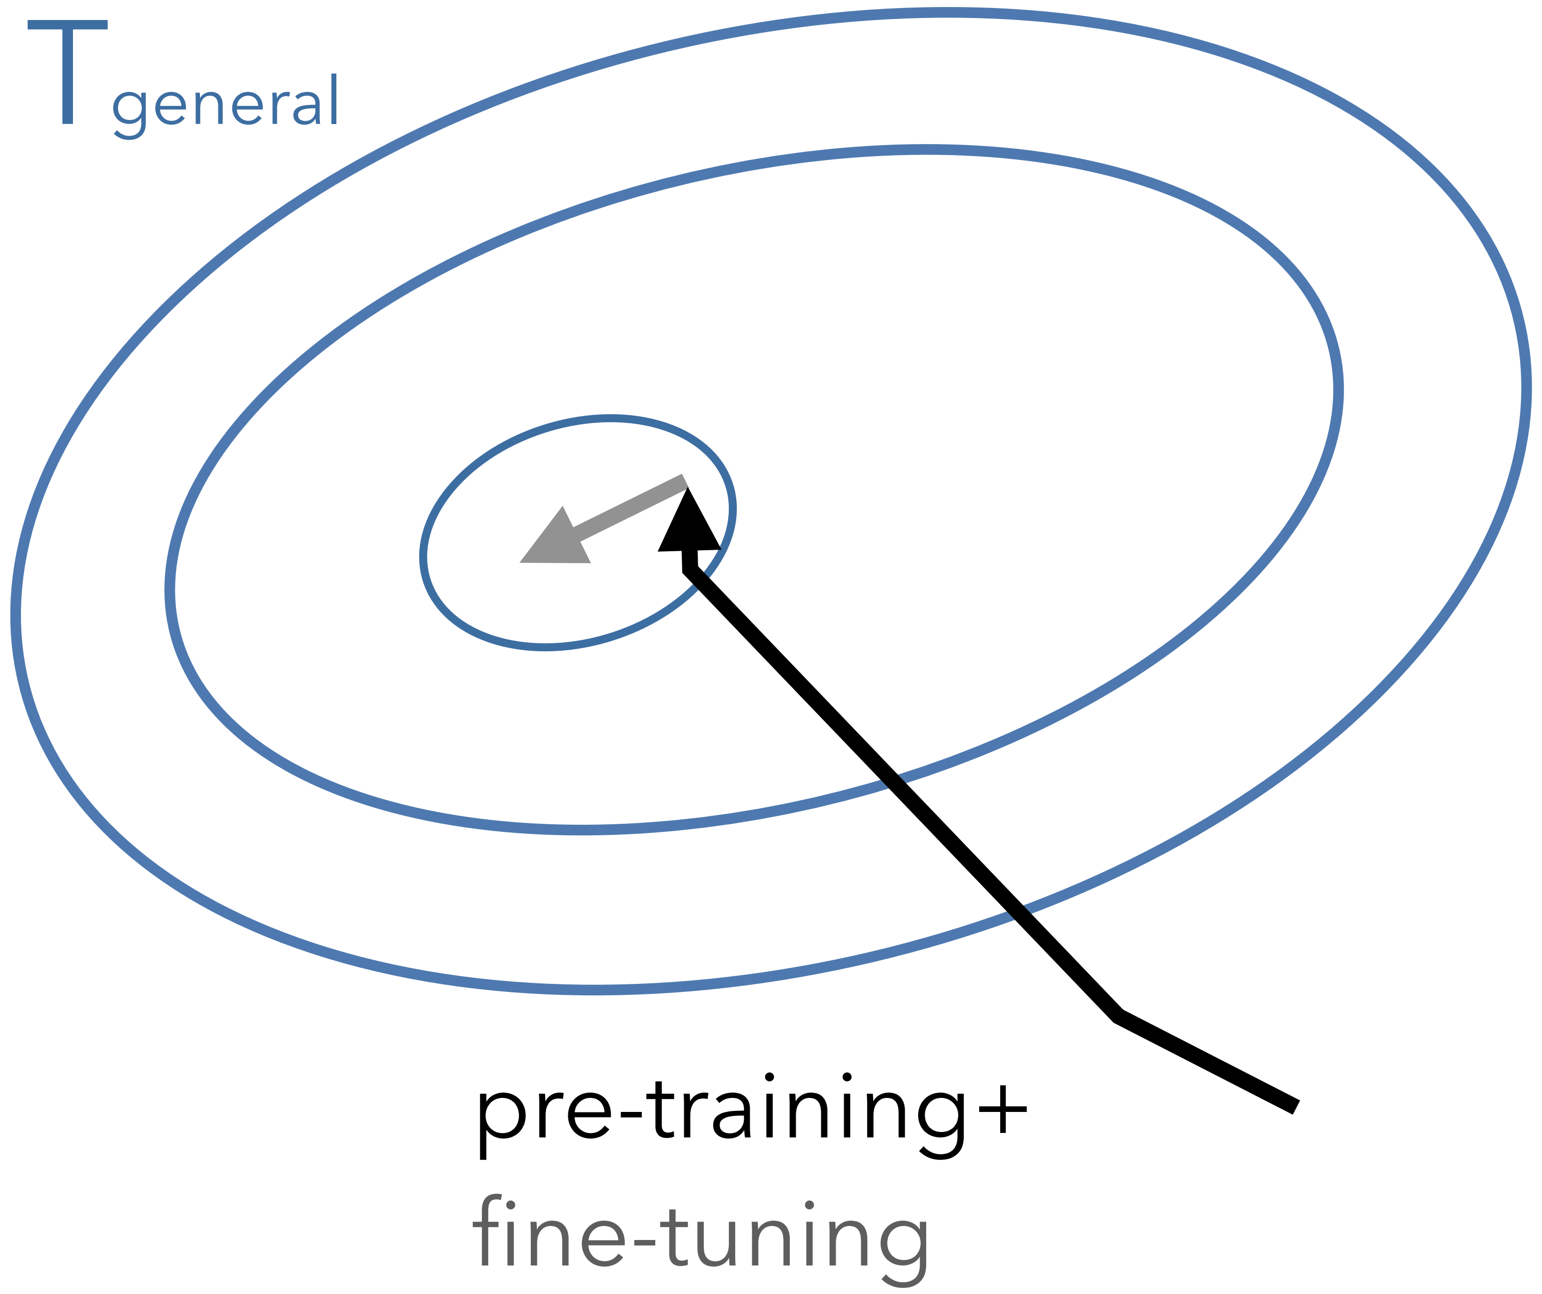

Domain-adaptive Pre-training

Gururangan et al. (2020)

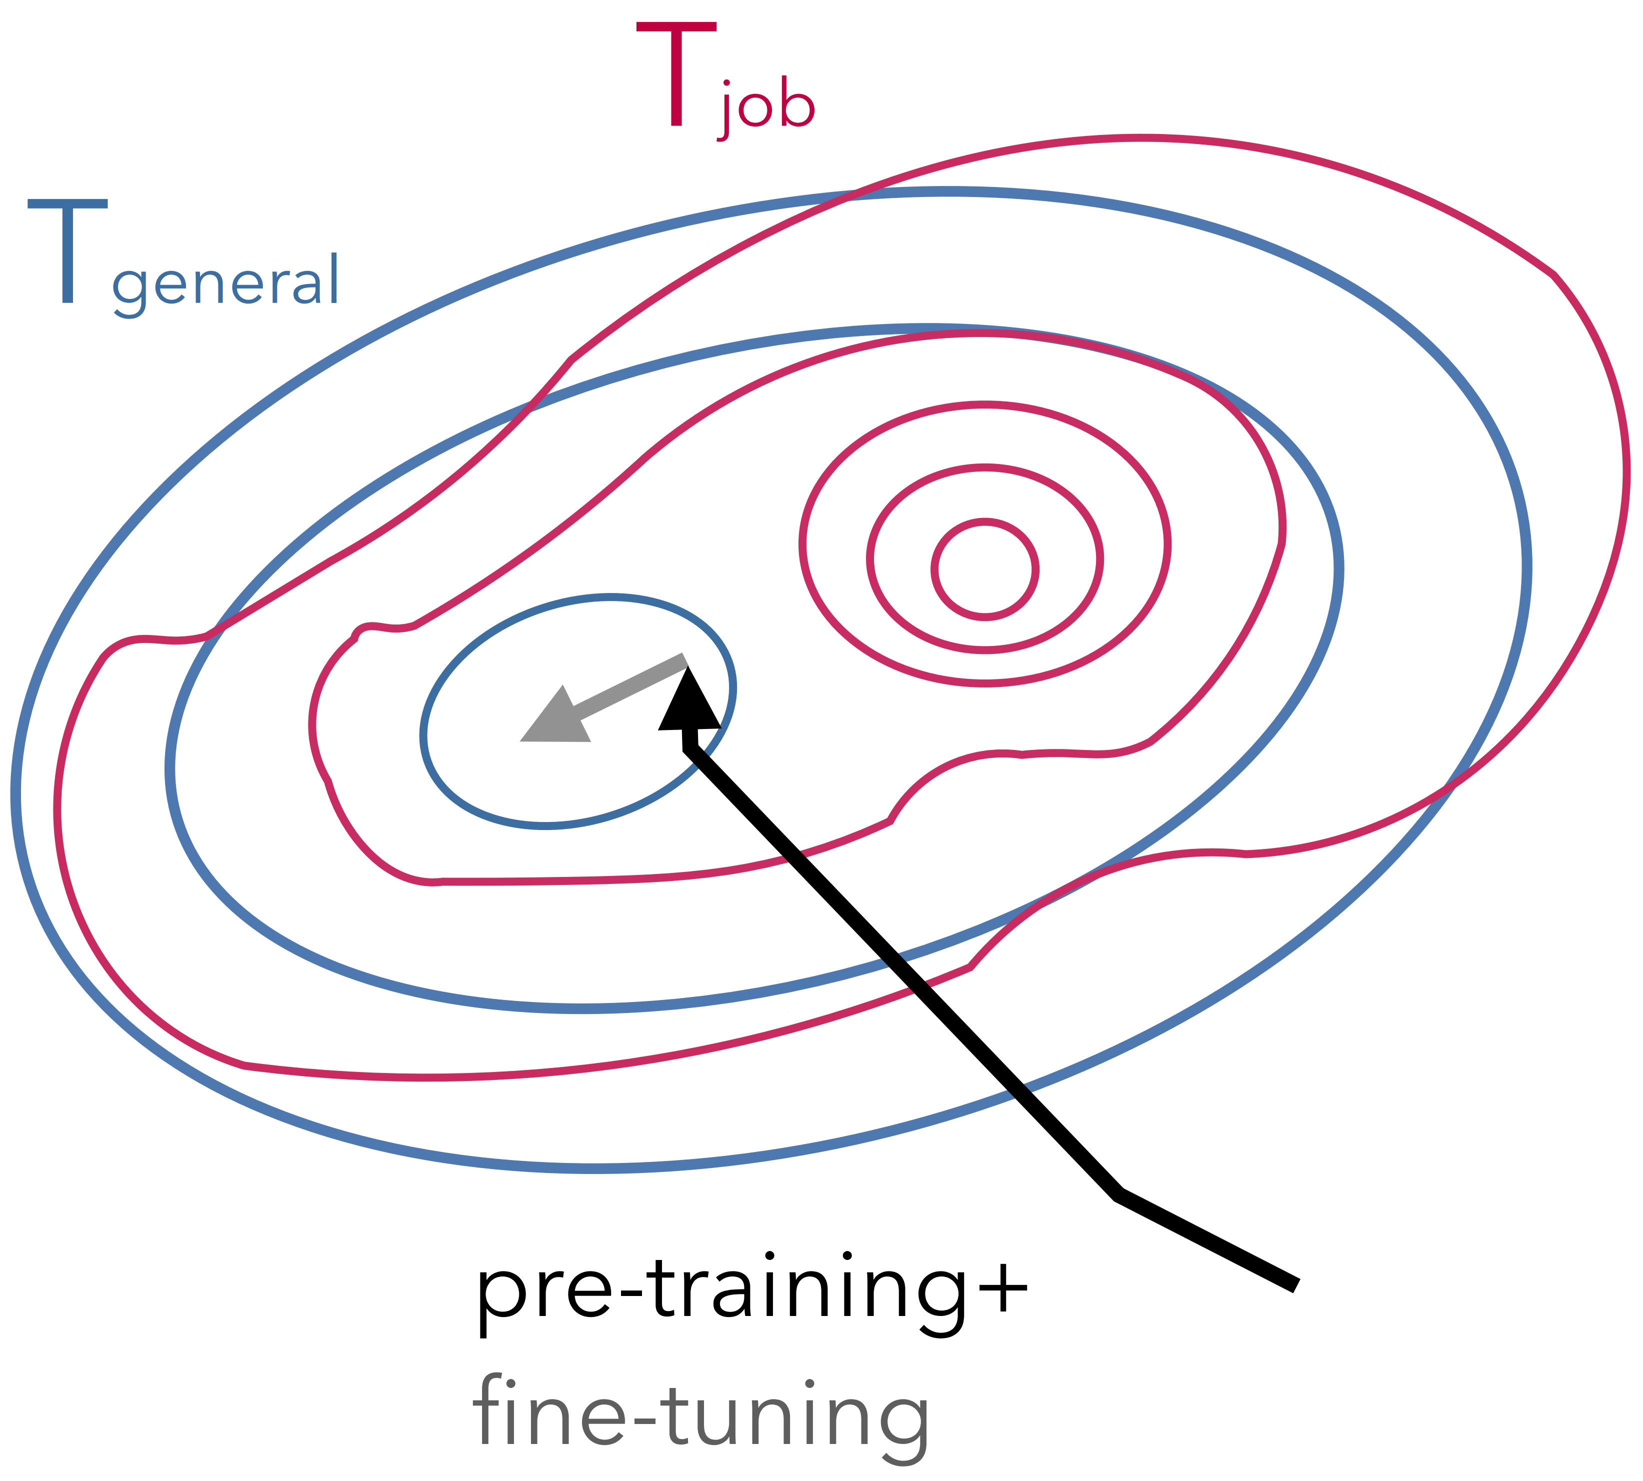

Domain-adaptive Pre-training

Gururangan et al. (2020)

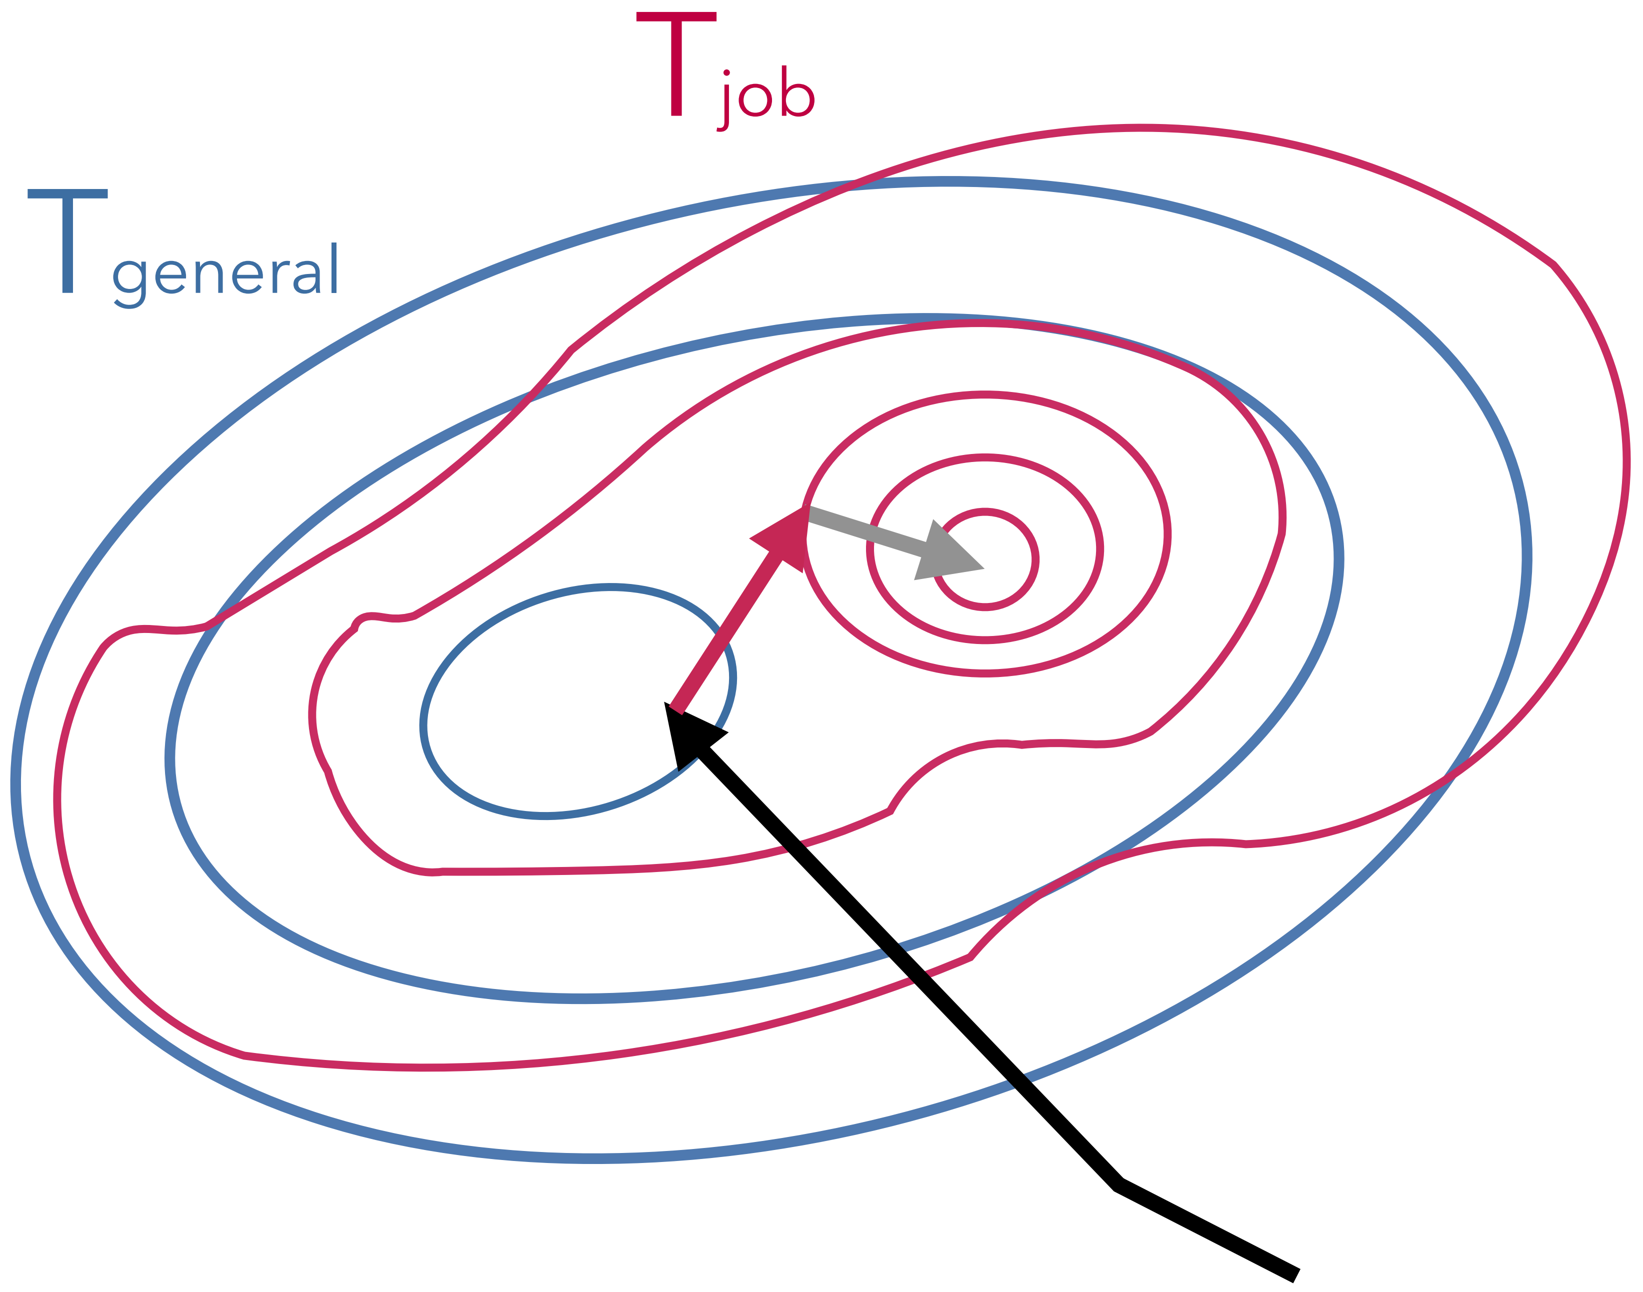

pre-training + domain-adaptive pre-training
+ fine-tuning

Single-task Learning

We can predict separately

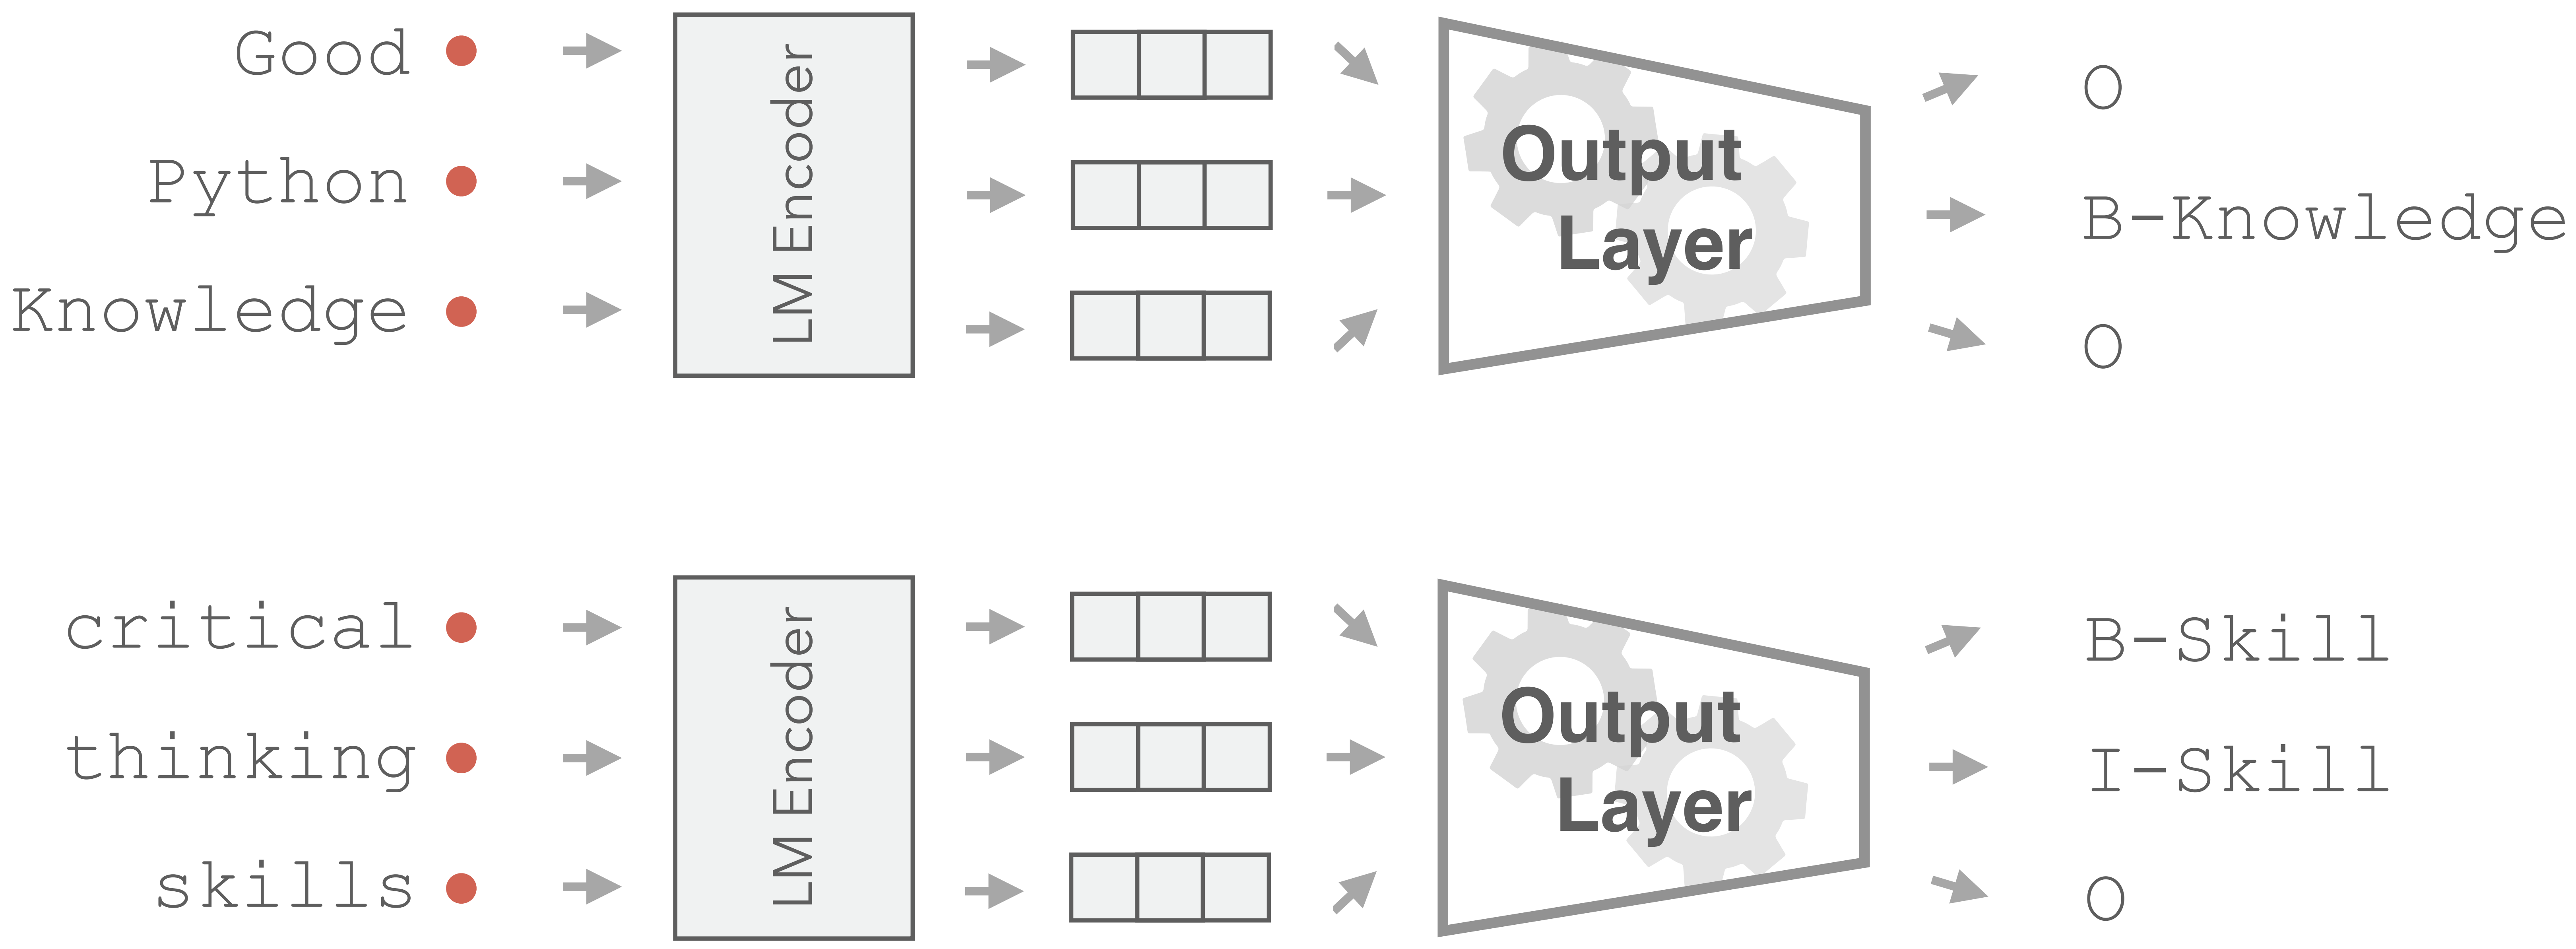

Multi-task Learning

Caruana (1997), hard parameter sharing

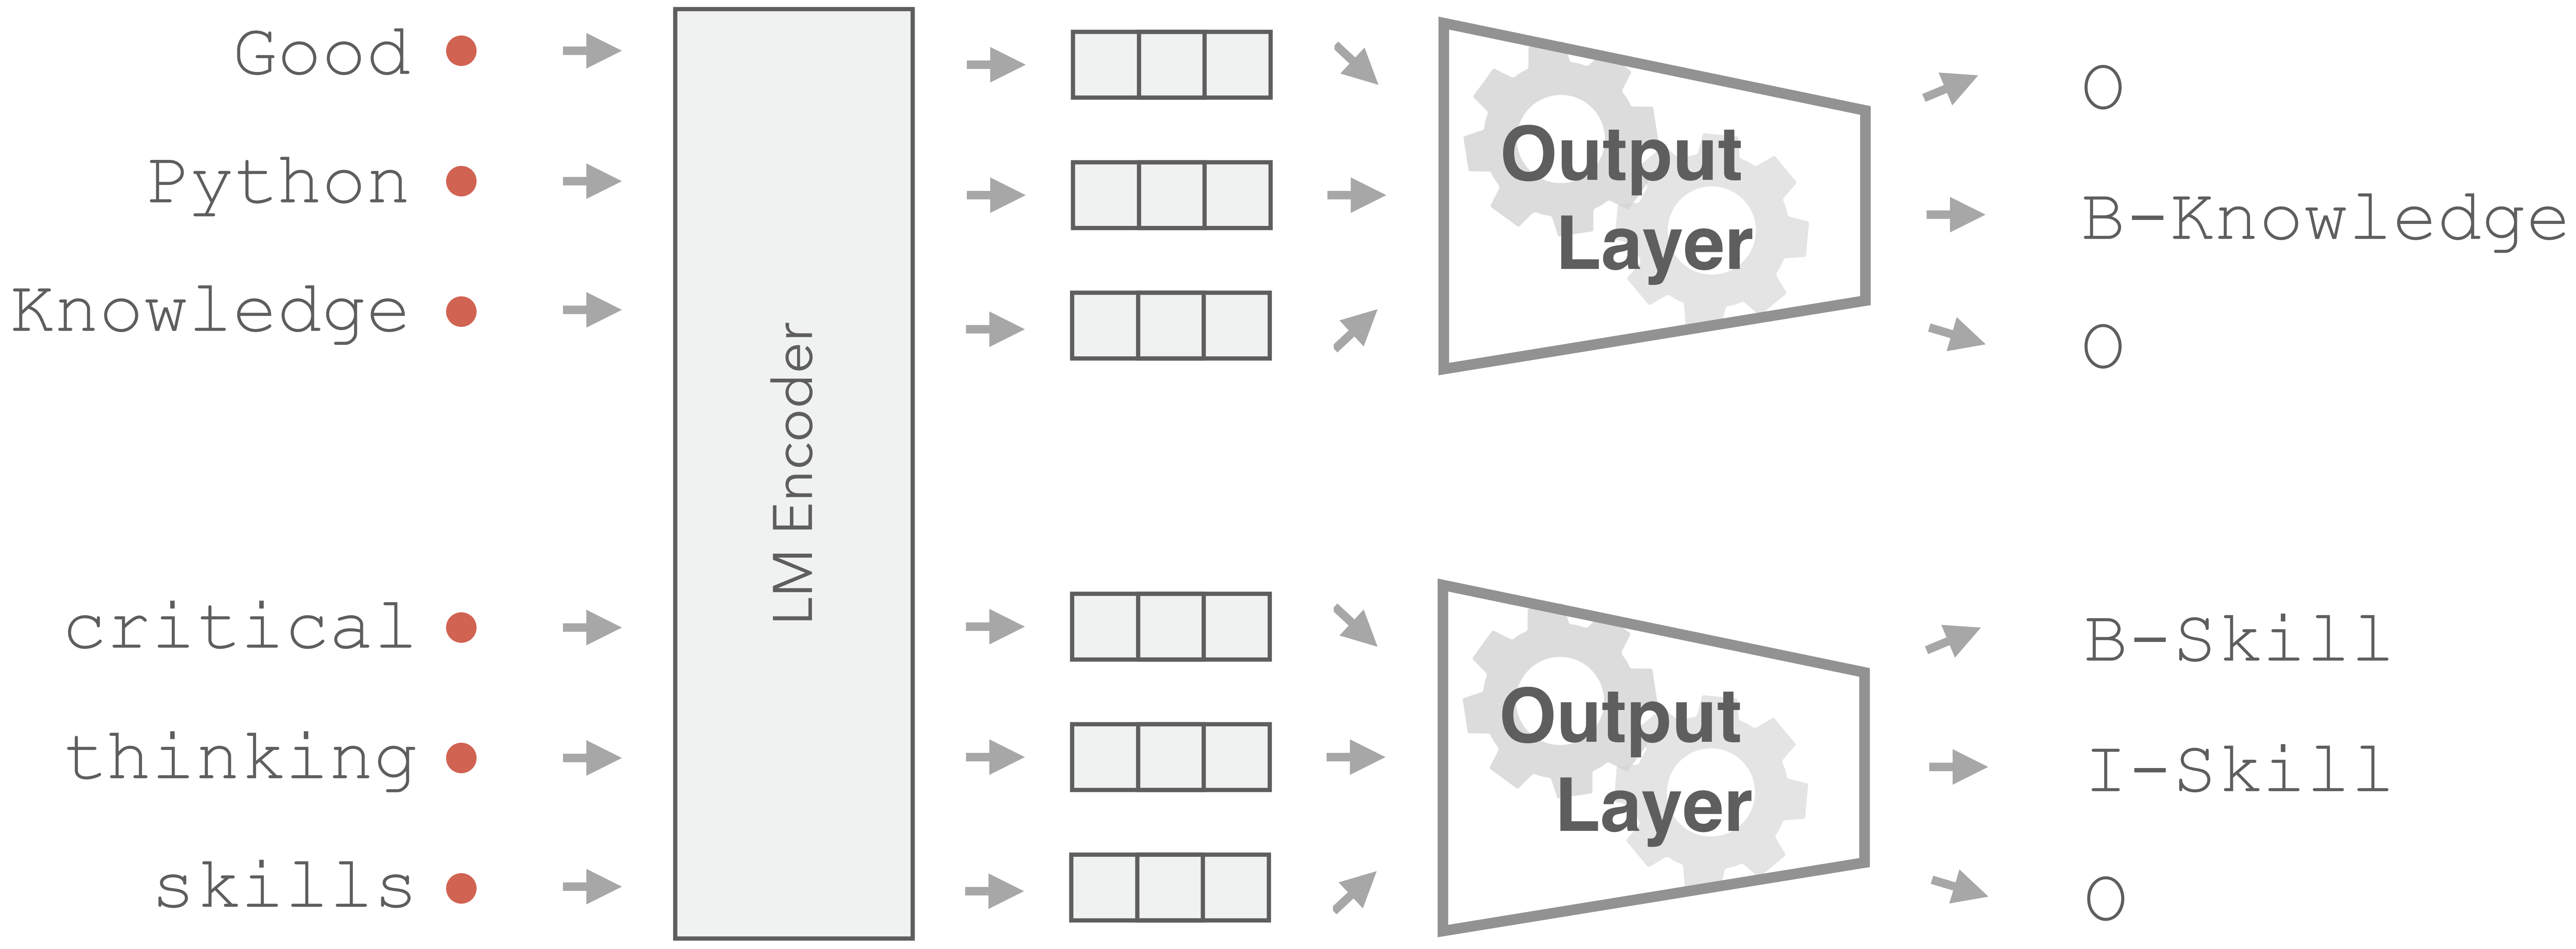

Setup and Results

Single Task vs. Multi Task Learning

- Split of **200/90/101** JPs.
- **Three** models:
  - **BERT** (Devlin et al., 2018)
  - **JobBERT** (ours, domain adapted)
  - **JobSpanBERT** (ours, domain adapted)
- Single-task vs. multi-task learning (skills, knowledge, or **combined**)
  - **Single-task performs better than multi-task learning in most cases**

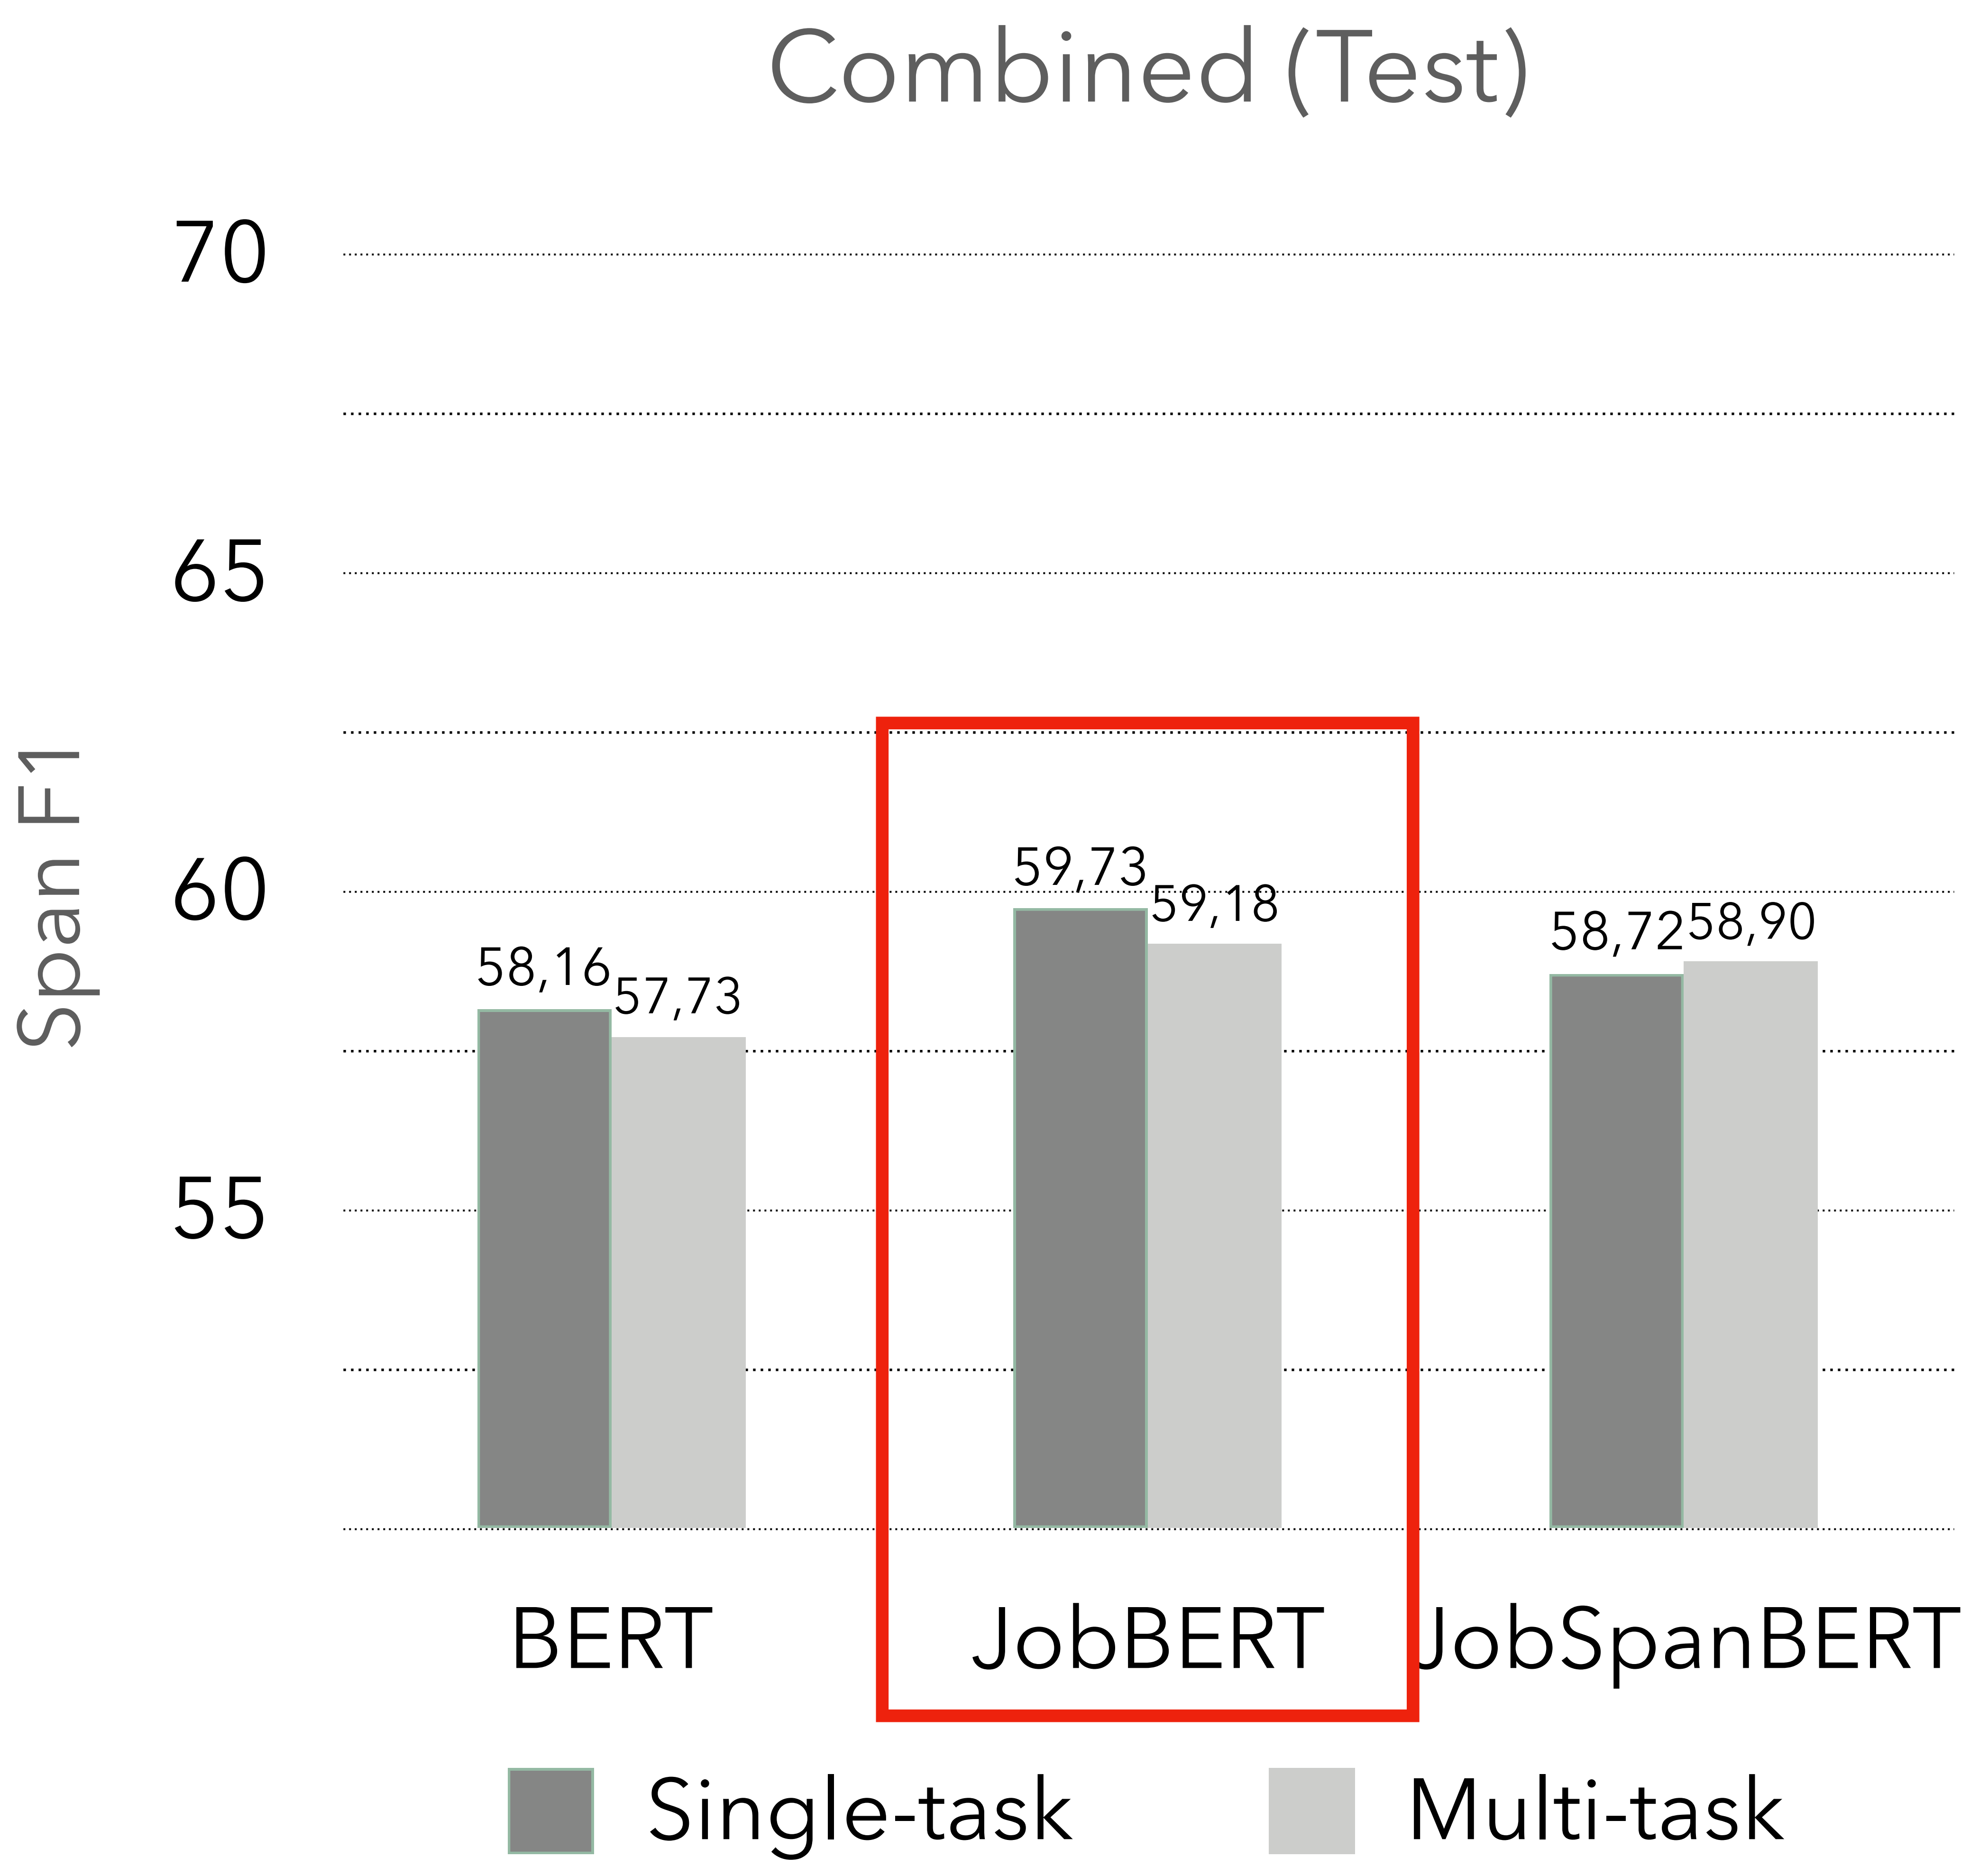

What does JobBERT work best on?

Challenge 2: Long spans

- JobBERT performs well on longer skills (3-7 tokens);
- Short skills include words such as “passionate” which could be a soft skill or not, could be challenging for the model;

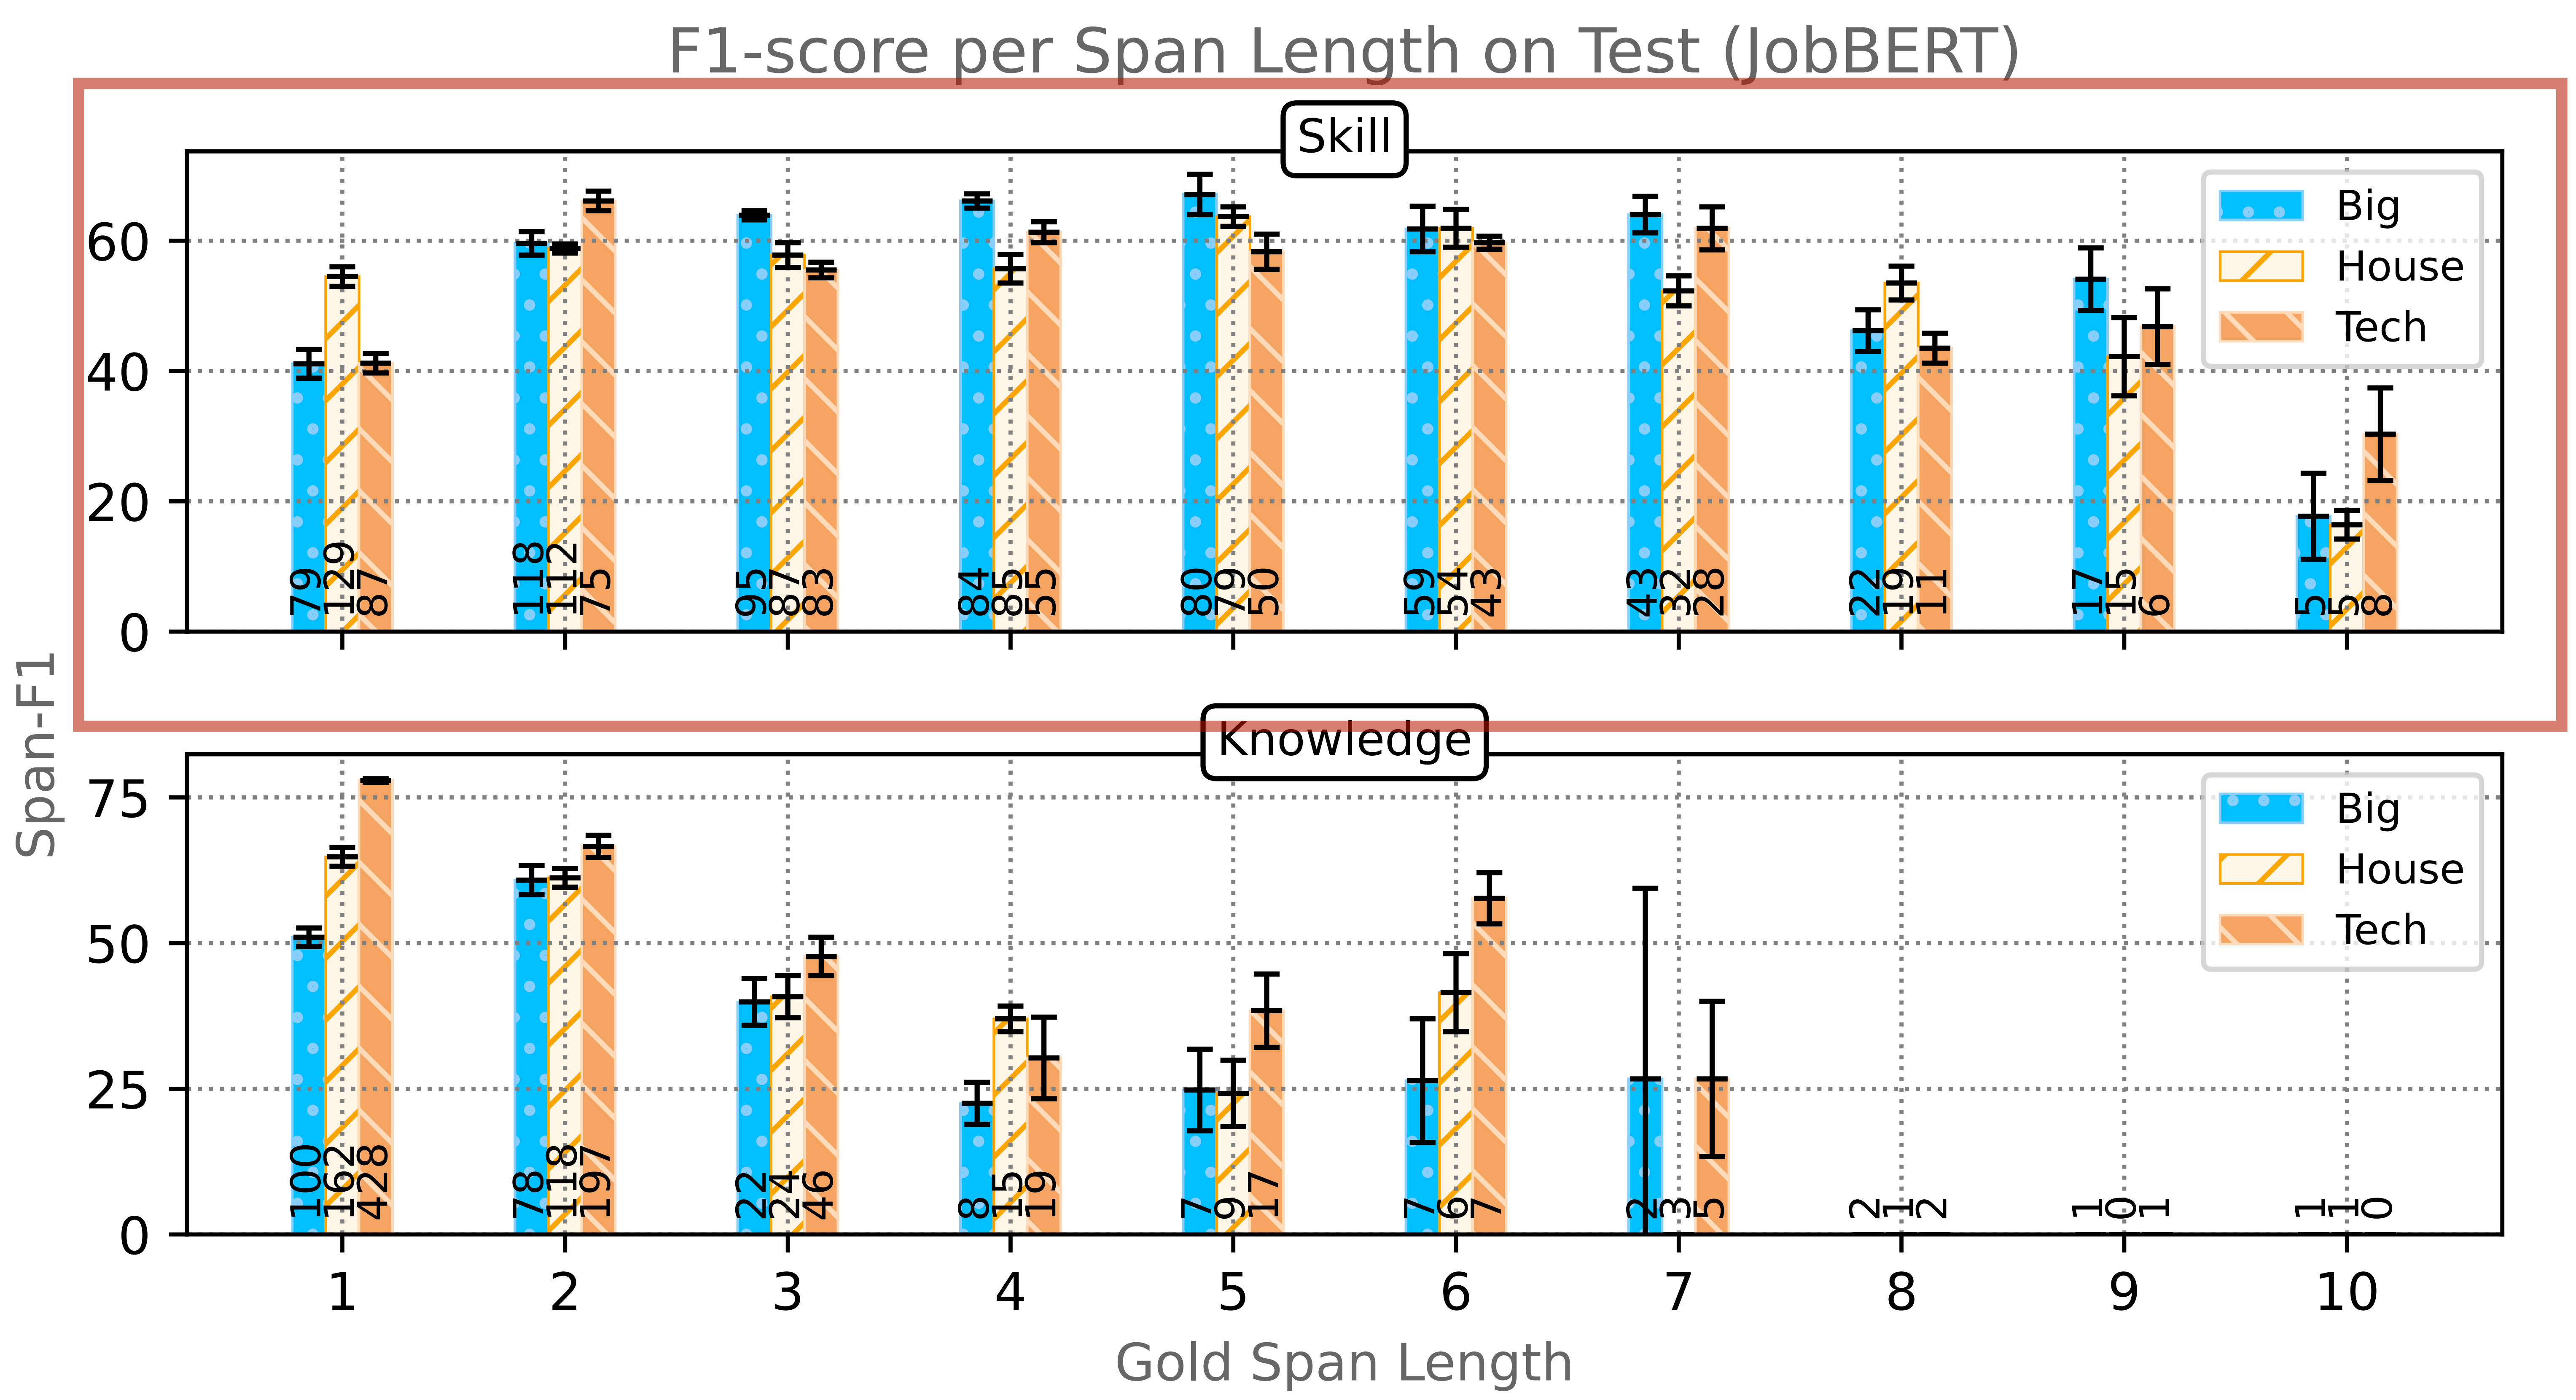

What does JobBERT work best on?

Challenge 2: Long spans

- JobBERT performs well on longer skills (3-7 tokens);
- Short skills include words such as “passionate” which could be a soft skill or not, could be challenging for the model;
- Knowledge components usually consist of proper nouns (e.g., Python, Java) and might be easier for the model to predict.

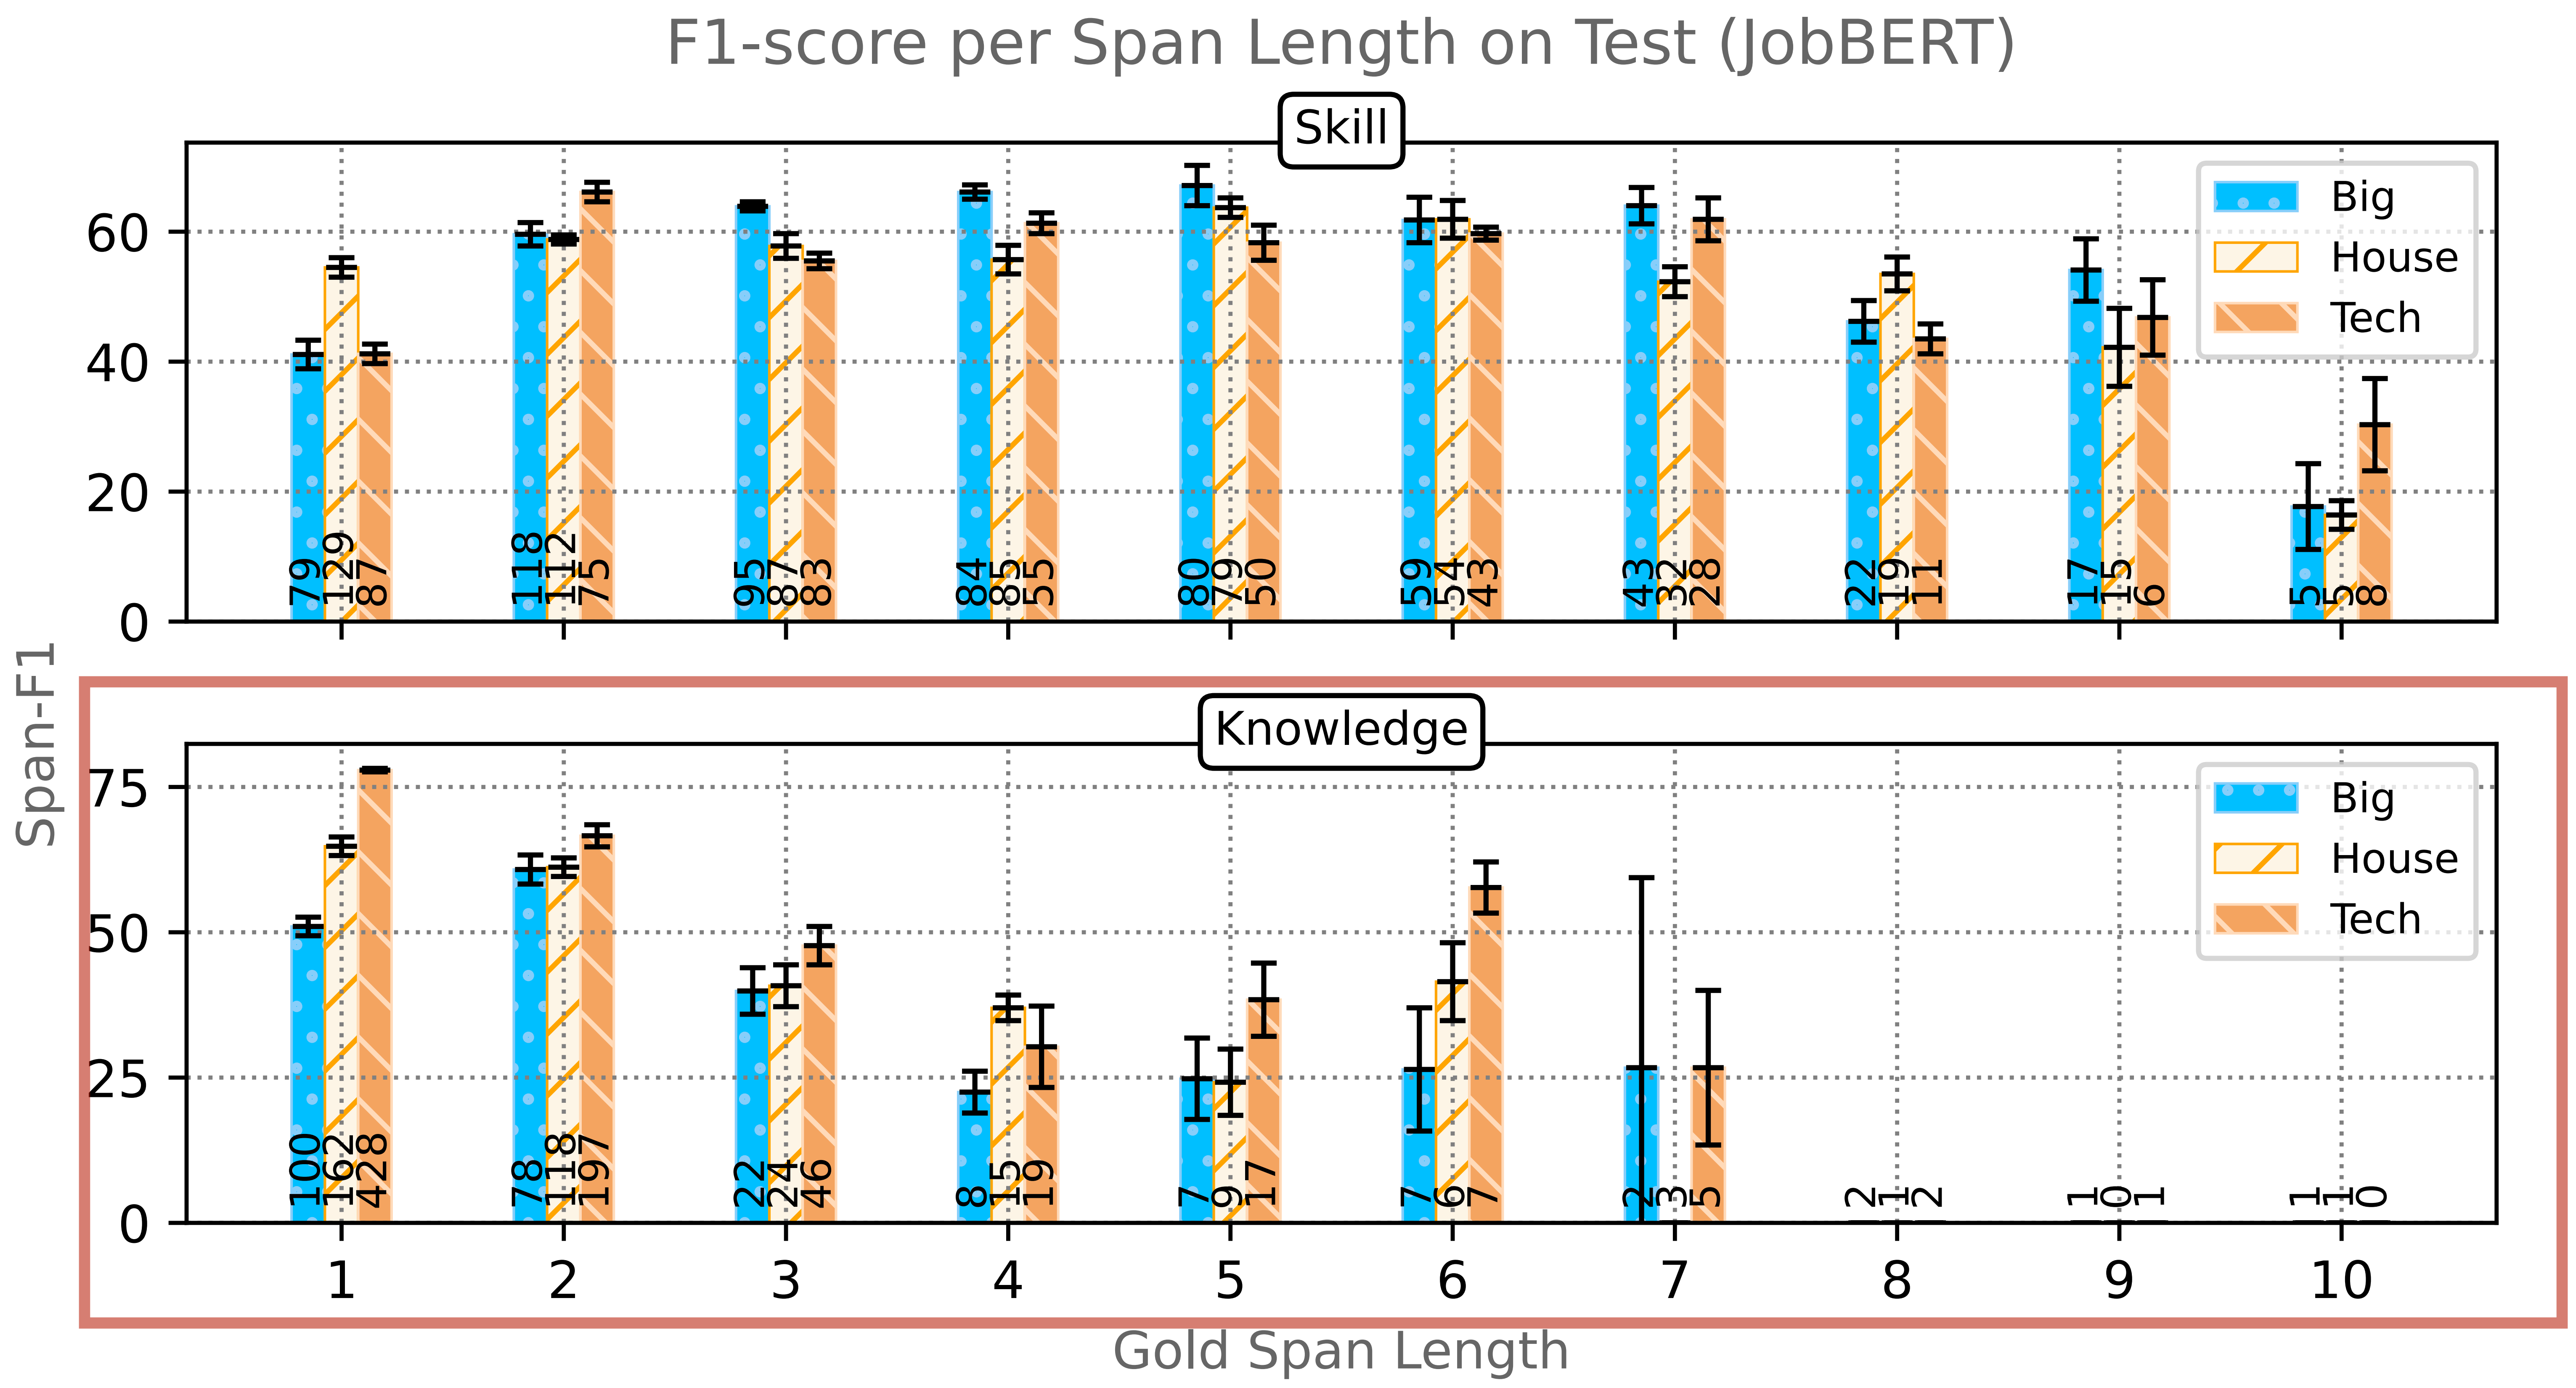

Takeaways

SkillSpan

- Domain-adaptive pre-training is a simple method to enhance your language model to a target domain.
- Knowledge components are easier to predict.
- Soft-like skills could pose a challenge to predict for language models.
- **Challenge 2, Span Length: Skills, by their nature, can exhibit considerable length, posing a challenge in extraction with NLP models.**
    - On skills, our model predicts long sequences fairly well (3-7 tokens), but starts to degrade in performance with longer spans.

What about other languages?

What about other languages?

Apart from English

- Not everyone is searching for jobs in English;
- How can we make extracted spans useful for stakeholders?
  - **Classifying** the extracted spans to a standardized label.
  - In turn, can be used for counting skill demand.

Kompetencer

 

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

Mike **Zhang**, Kristian Nørgaard **Jensen**, and Barbara **Plank**.

In Proceedings of LREC 2022.

Similar to SkillSpan

But with extensions.

- Similar to SkillSpan, we annotated for skill and knowledge components in job postings.
- **Two extensions**:
  - Danish;
  - Classifying spans for their ESCO taxonomy code.

How did you get the annotations?

Challenge 5: Output Space.

- We rely on distant supervision (Mintz et al. 2009)
- Query ESCO API with the span to (hopefully) get a taxonomy code out.
  - In this work, we only consider the first major group (26 classes).

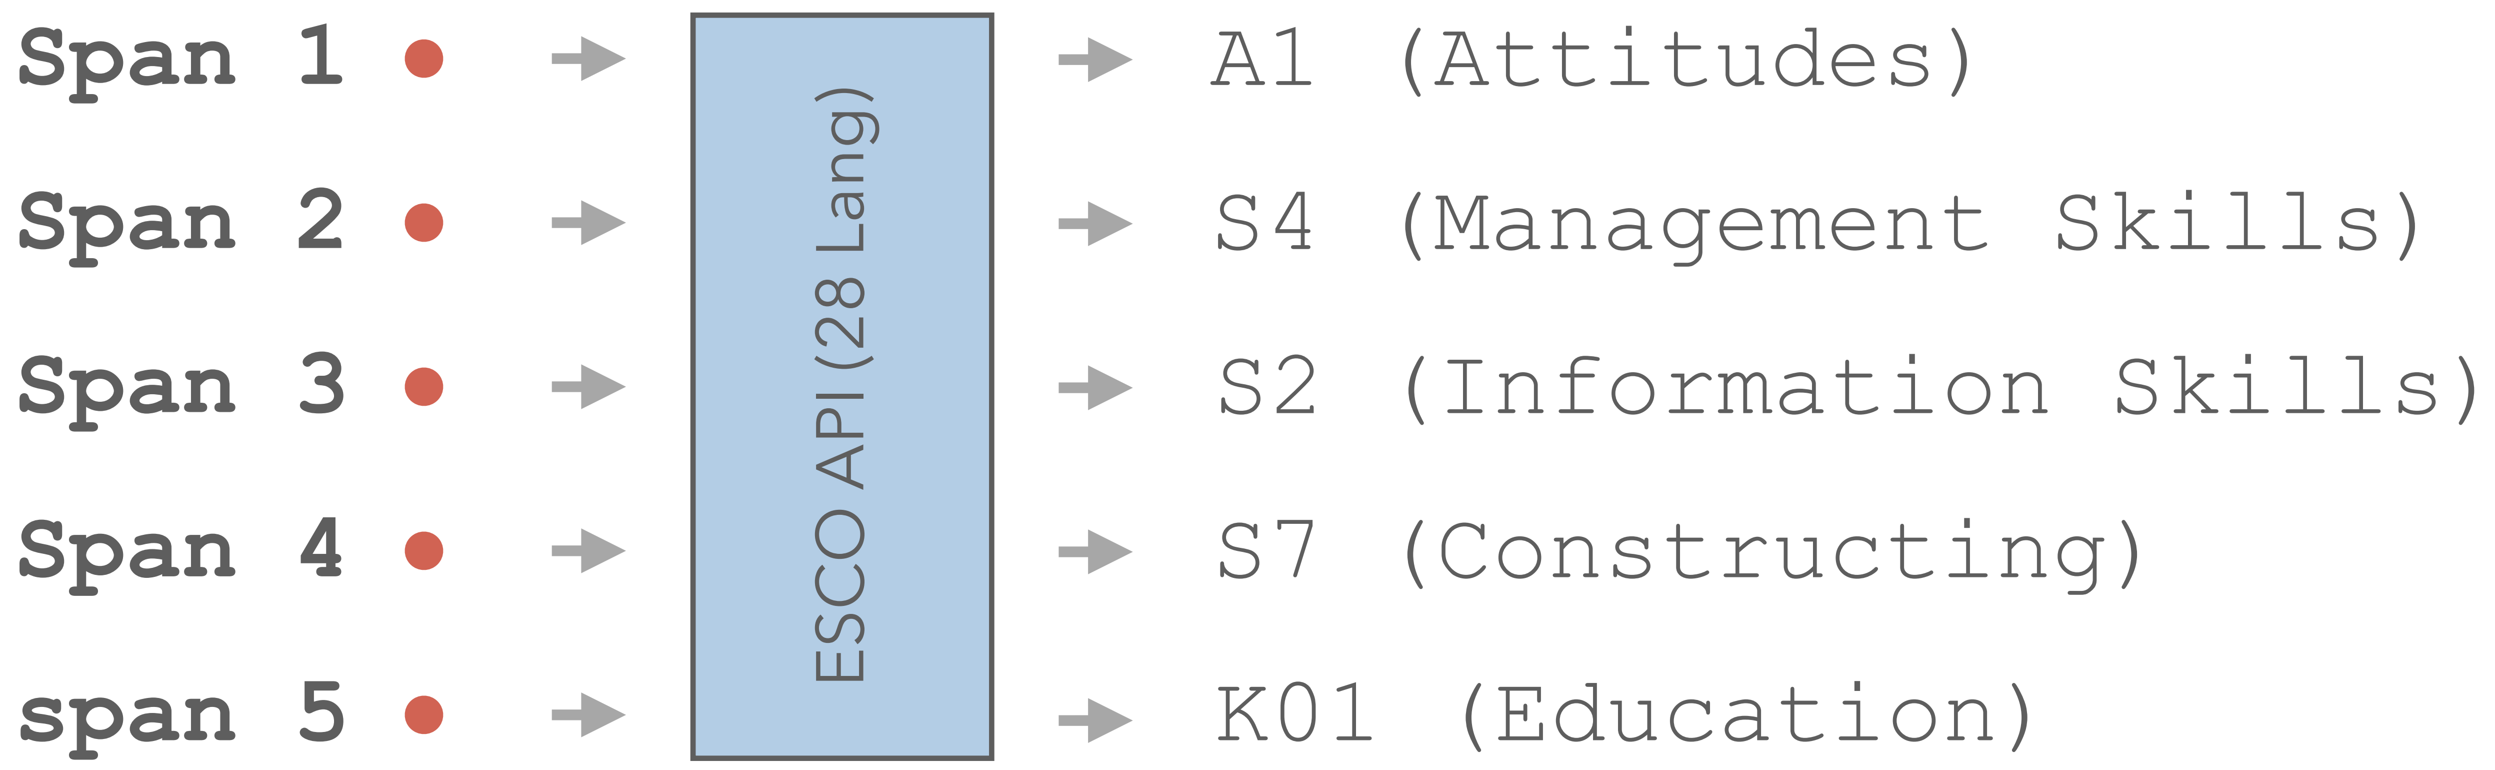

Setting Context

Classifying Spans == Multiclass Classification (in this case).

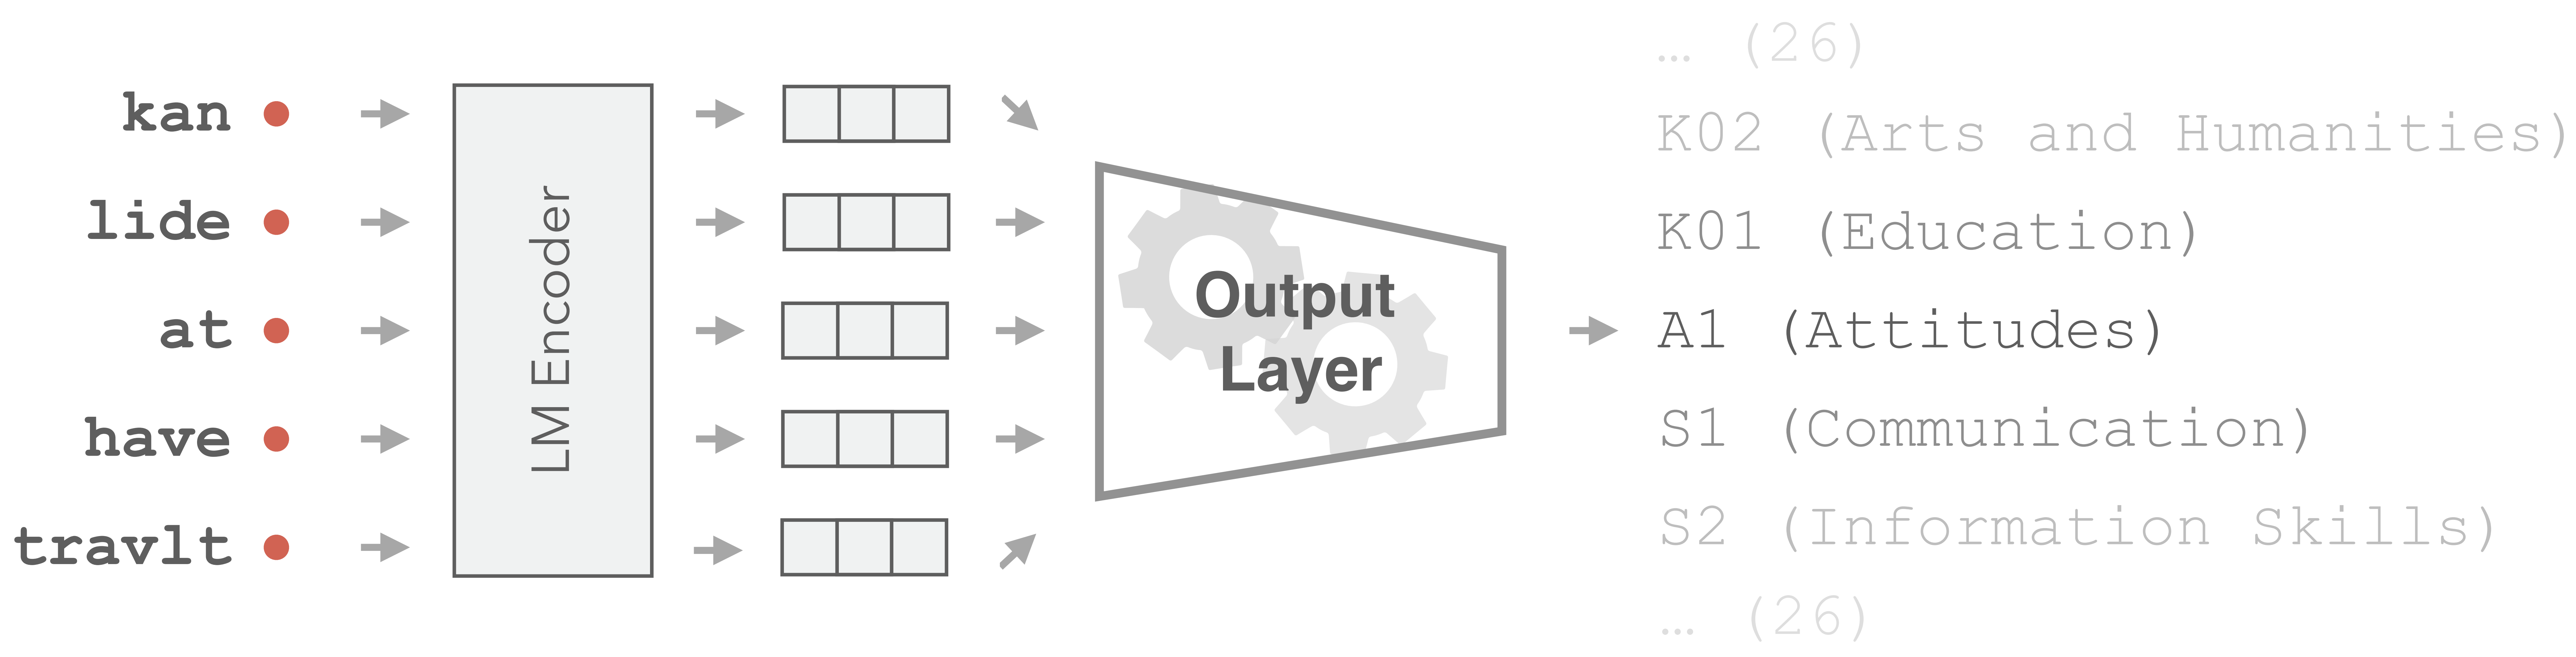

"Individual work styles that can affect how well someone performs a job."

Key Statistics

Dataset statistics.

- 60 job postings
- 1.5K sentences
  - 20.3K tokens
- ~900 annotations
- 26 classes
- Data:
  - Skillspan (EN) and;
  - Kompetencer (DA)

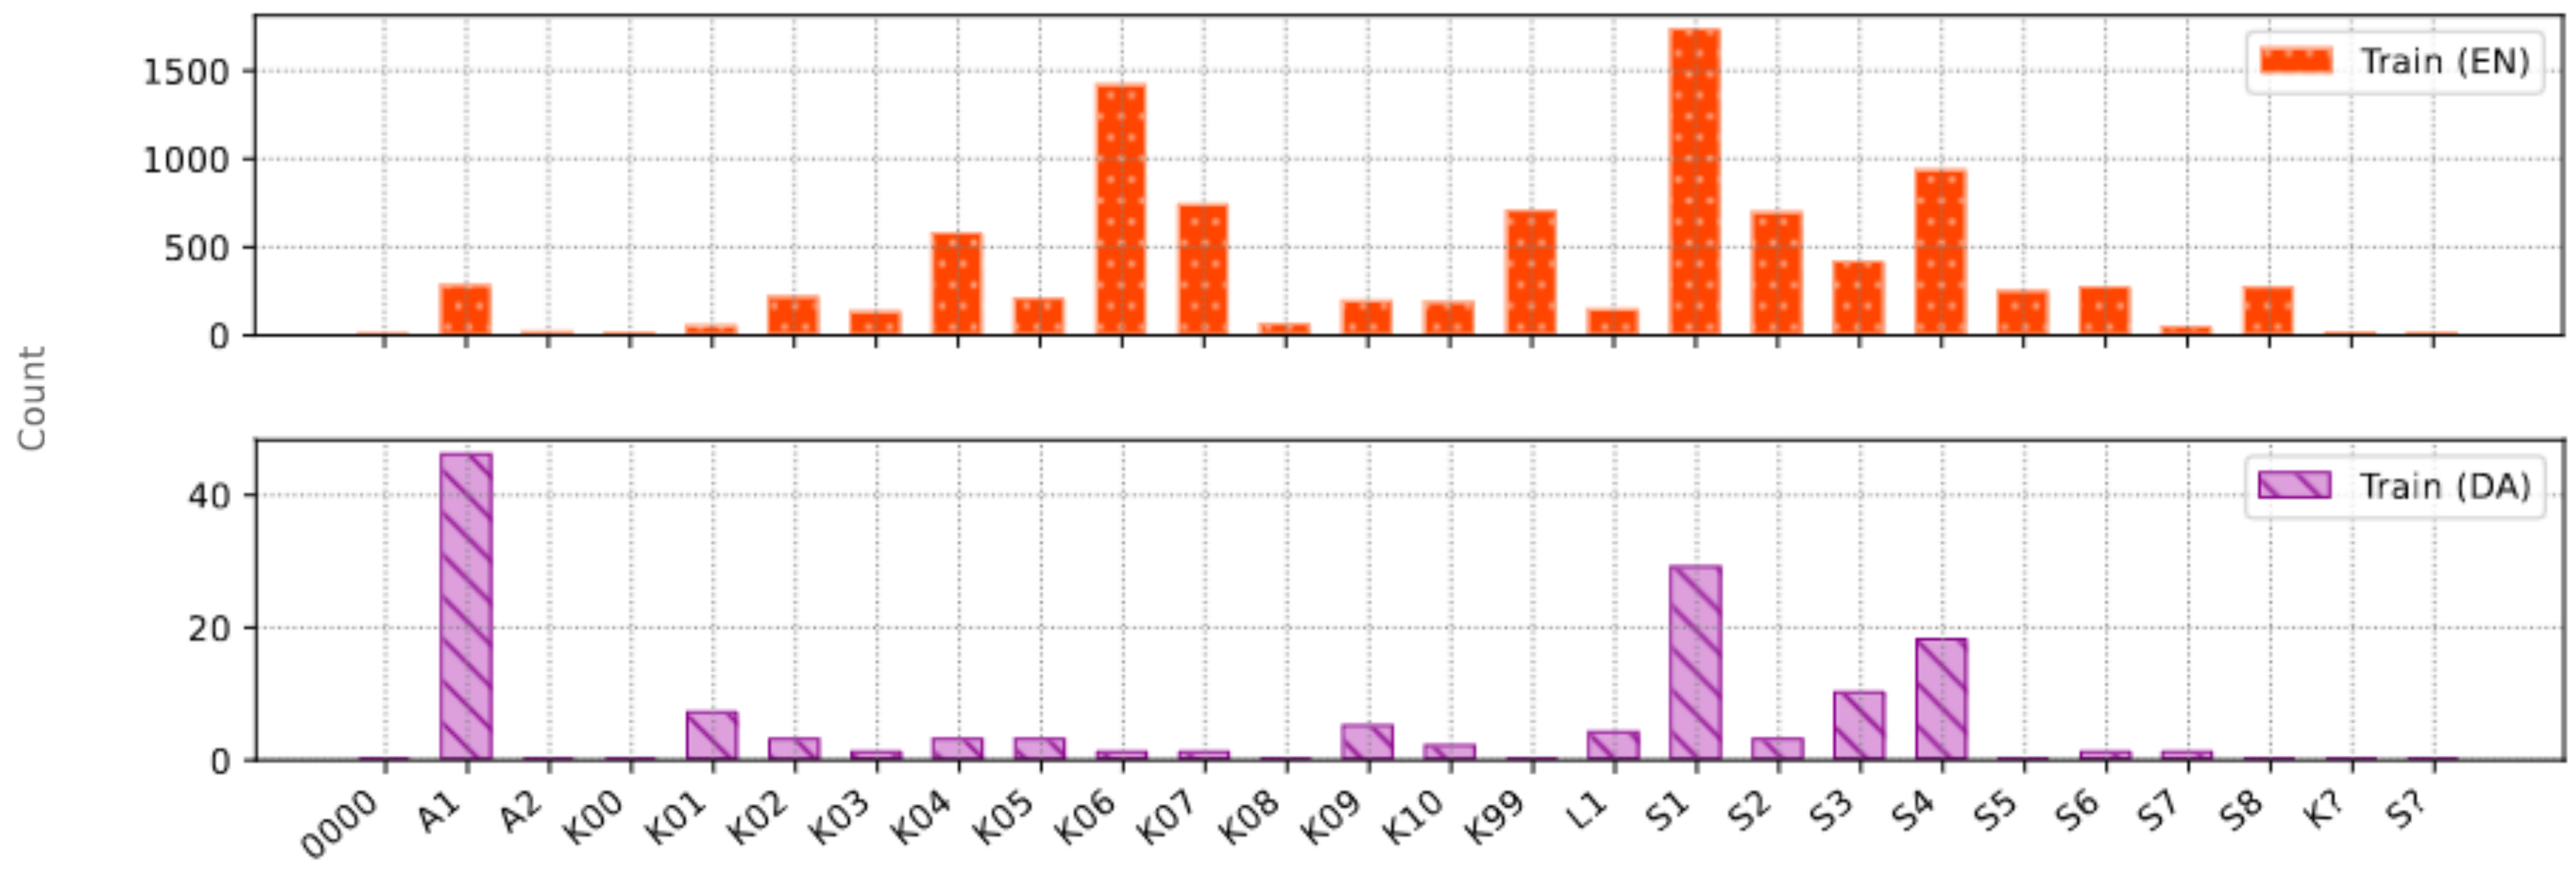

Setup and Results

Monolingual vs. Multilingual Models

- Split of **10 train (few-shot) / 50 test DA** JPs.
- **Zero-shot** (train on EN, apply to EN/DA)

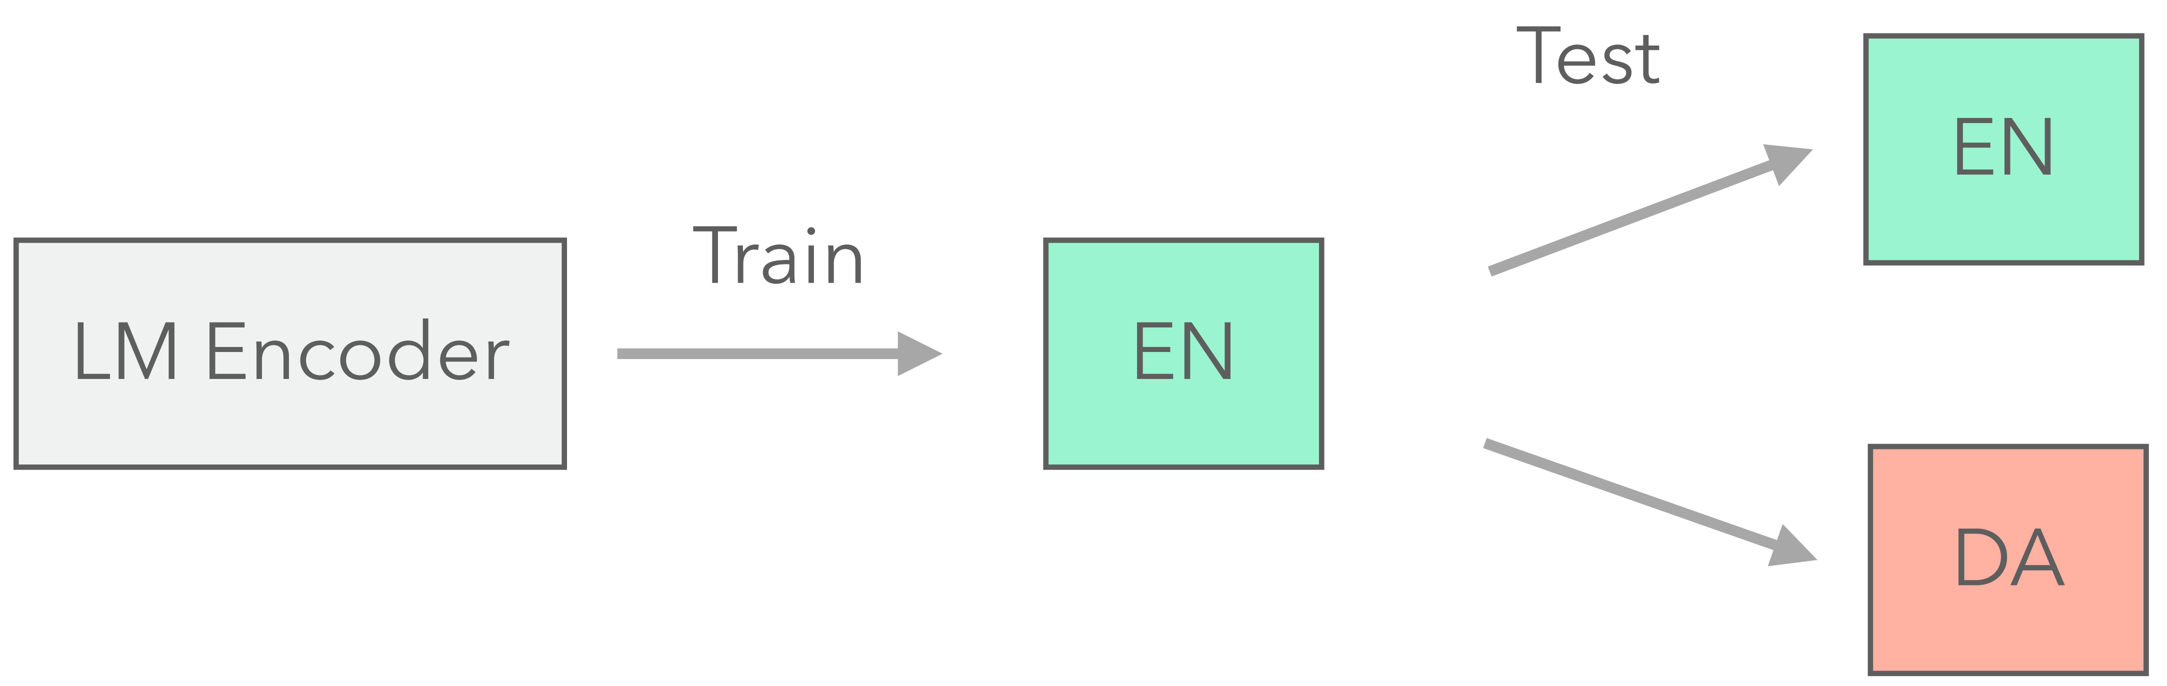

Setup and Results

Monolingual vs. Multilingual Models

- Split of **10 train (few-shot) / 50 test DA** JPs.
- **Zero-shot** (train on EN, apply to EN/DA)
- **Few-shot** (train on 10 DA JPs, apply to EN/DA)

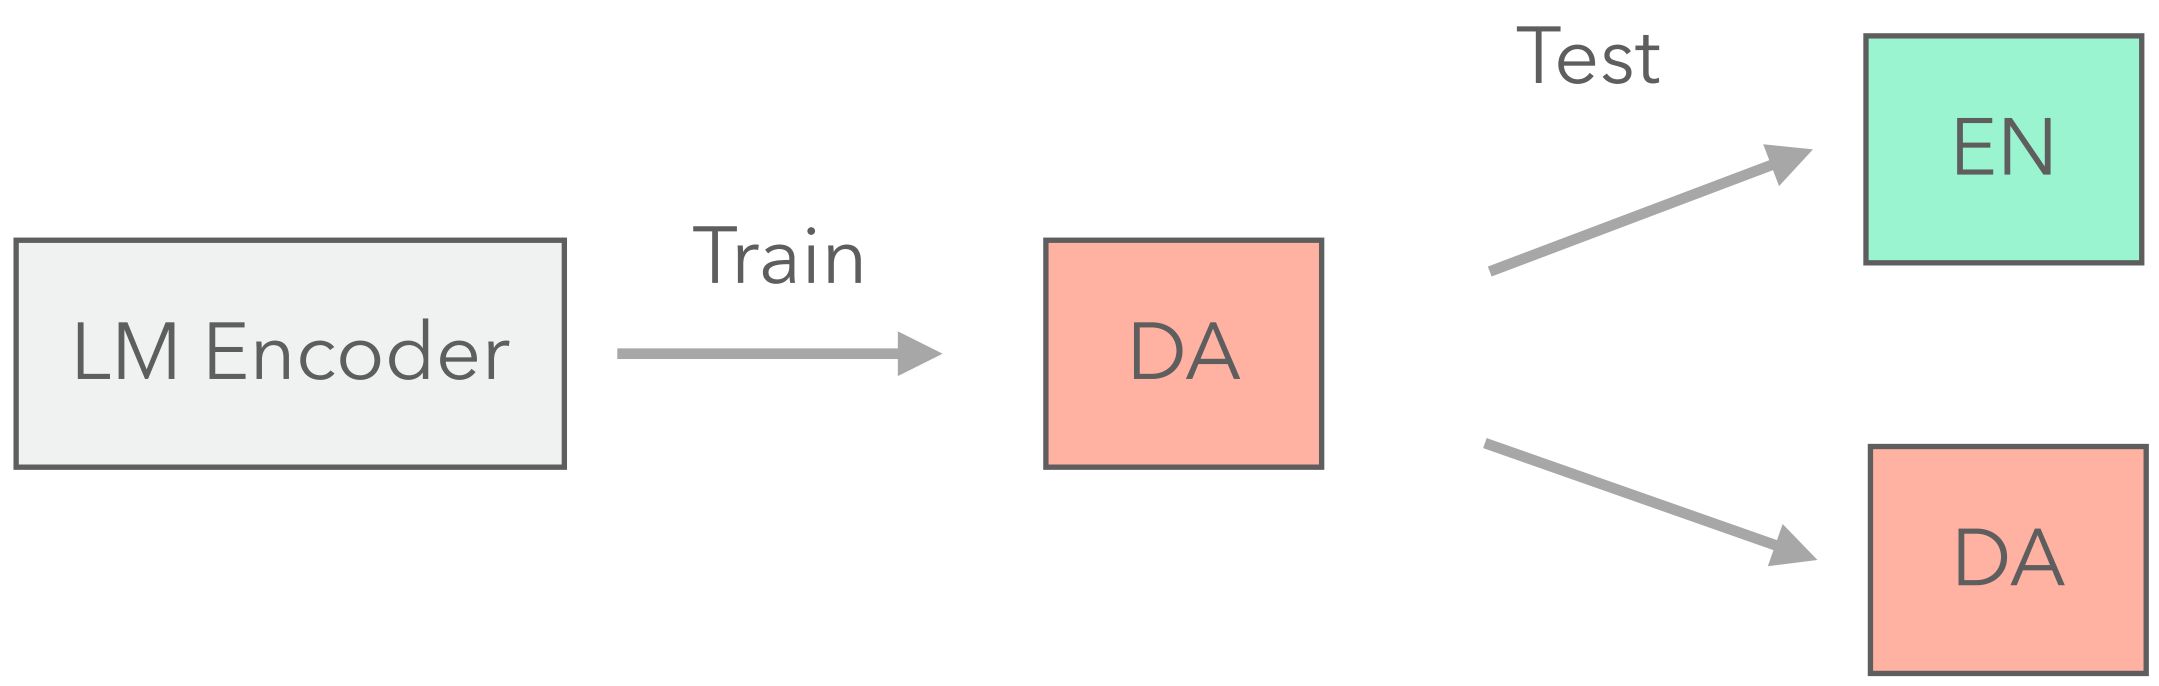

Setup and Results

Monolingual vs. Multilingual Models

- Split of **10 train (few-shot) / 50 test DA** JPs.
- **Zero-shot** (train on EN, apply to EN/DA), **Few-shot** (train on 10 DA JPs, apply to EN/DA)
- **Five** models:
  - EN: **BERT** (Devlin et al., 2018), **JobBERT** (Zhang et al., 2022)
  - DA: **DaBERT** (botxo), **DaJobBERT** (ours) (again domain-adaptive pre-training)
  - Multilingual: **RemBERT** (Chung et al., 2020)

Setup and Results

Monolingual vs. Multilingual Models

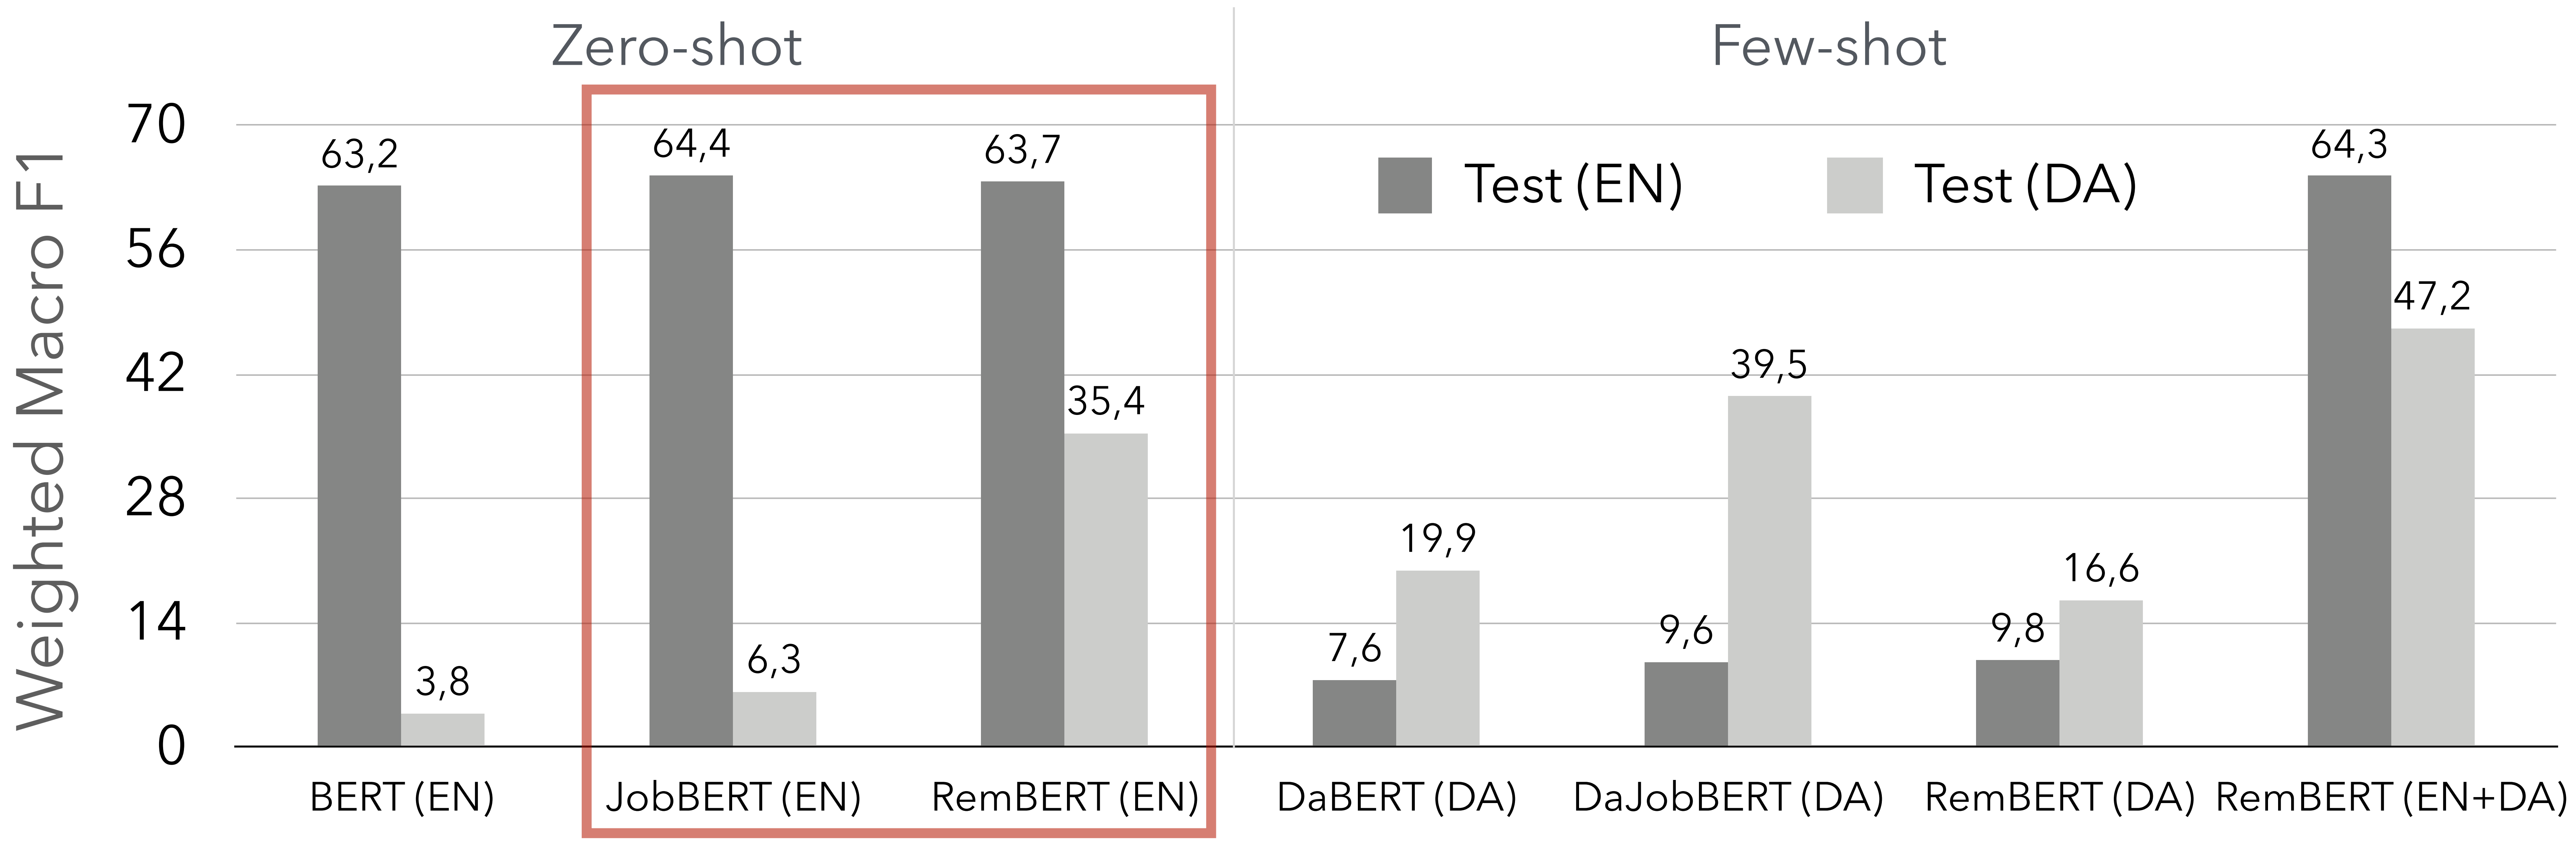

Setup and Results

Monolingual vs. Multilingual Models

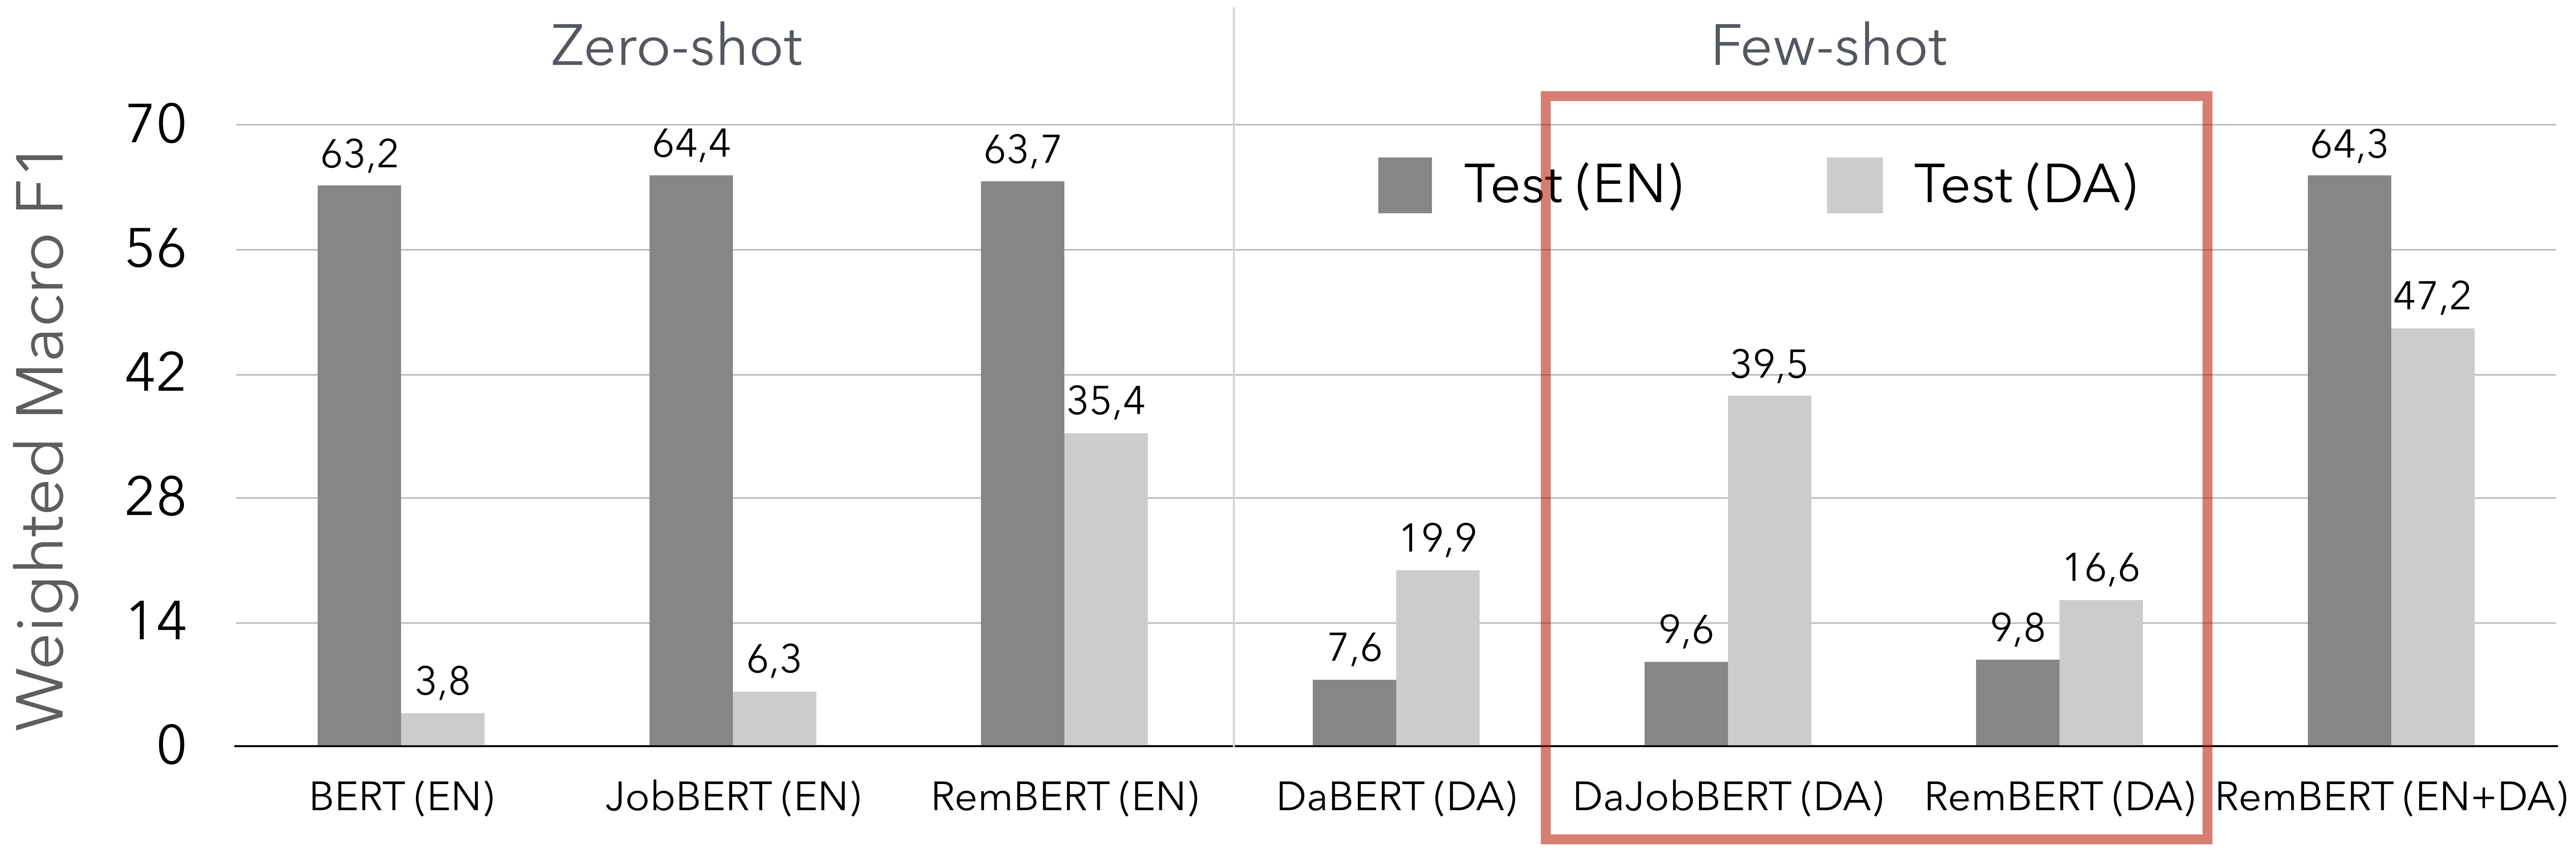

Setup and Results

Monolingual vs. Multilingual Models

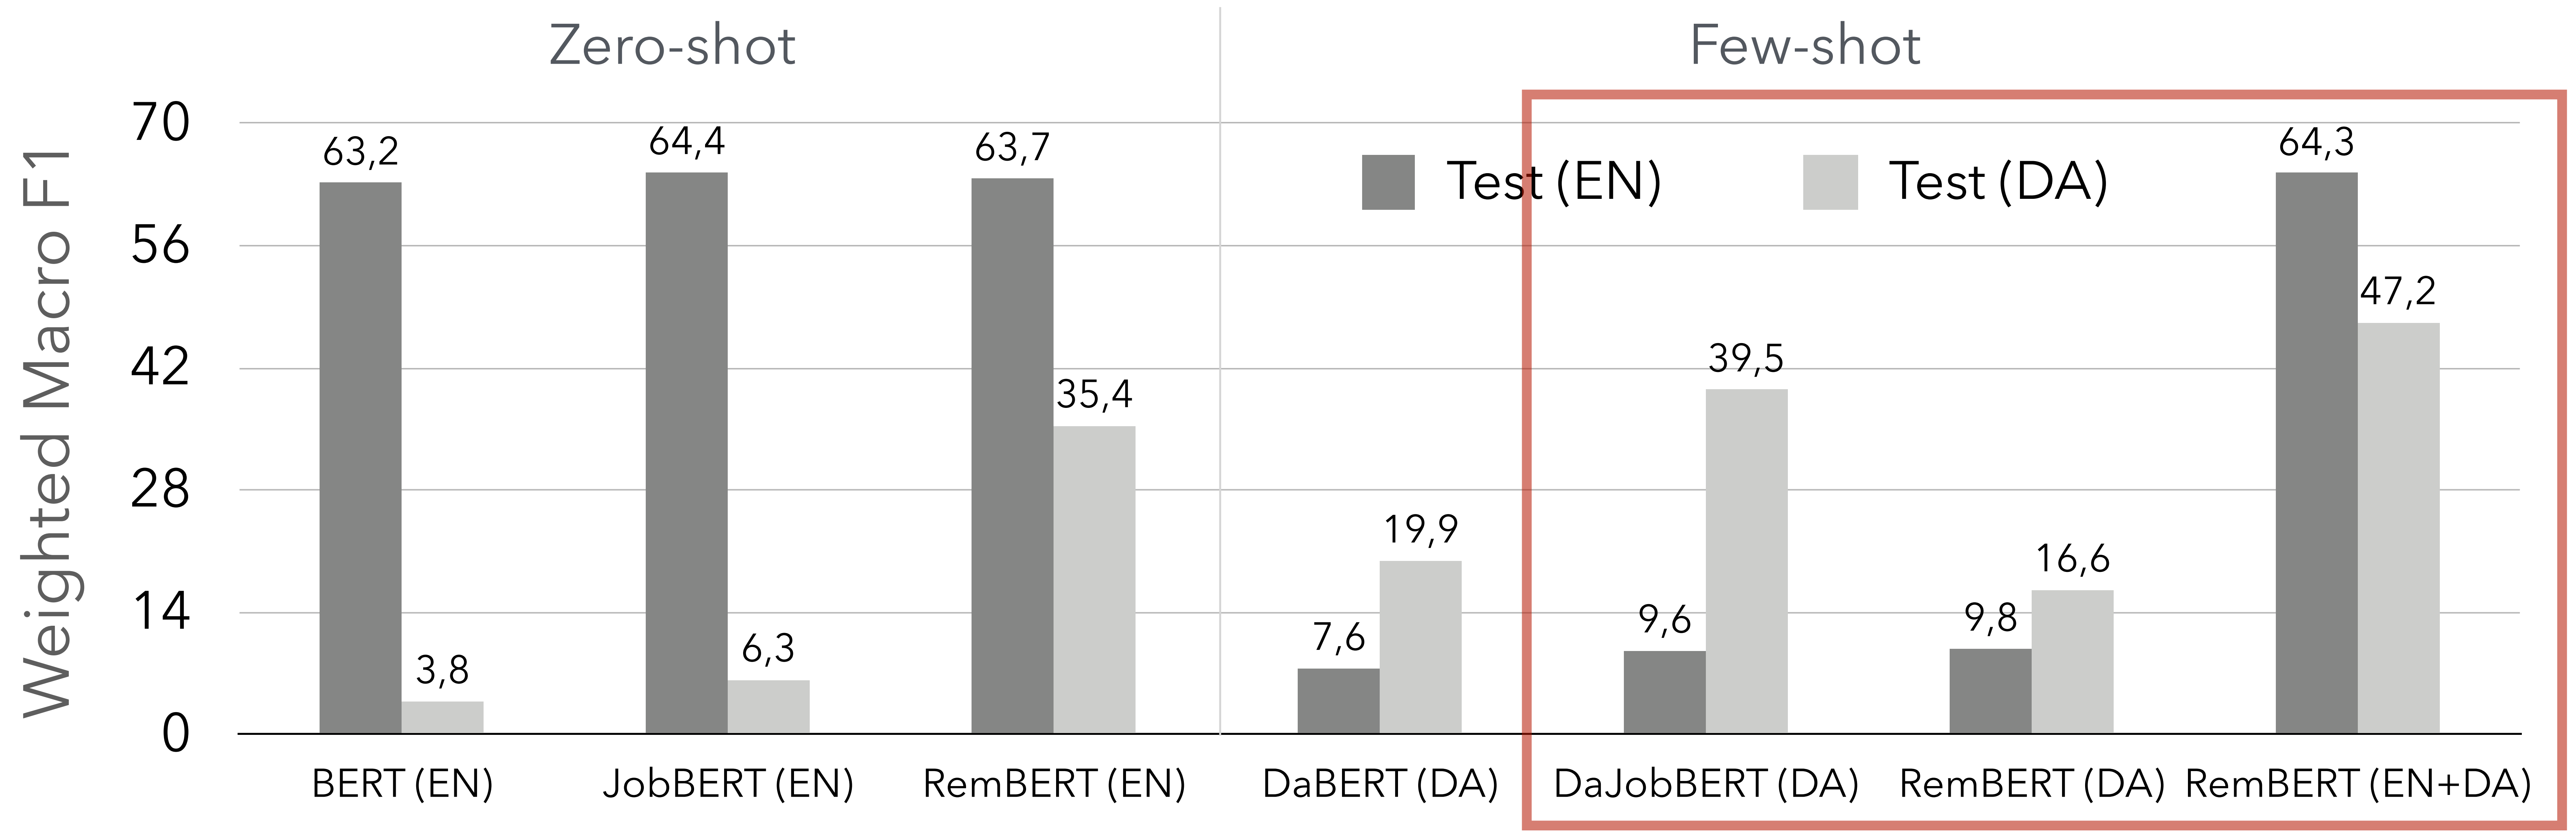

Setup and Results

Monolingual vs. Multilingual Models

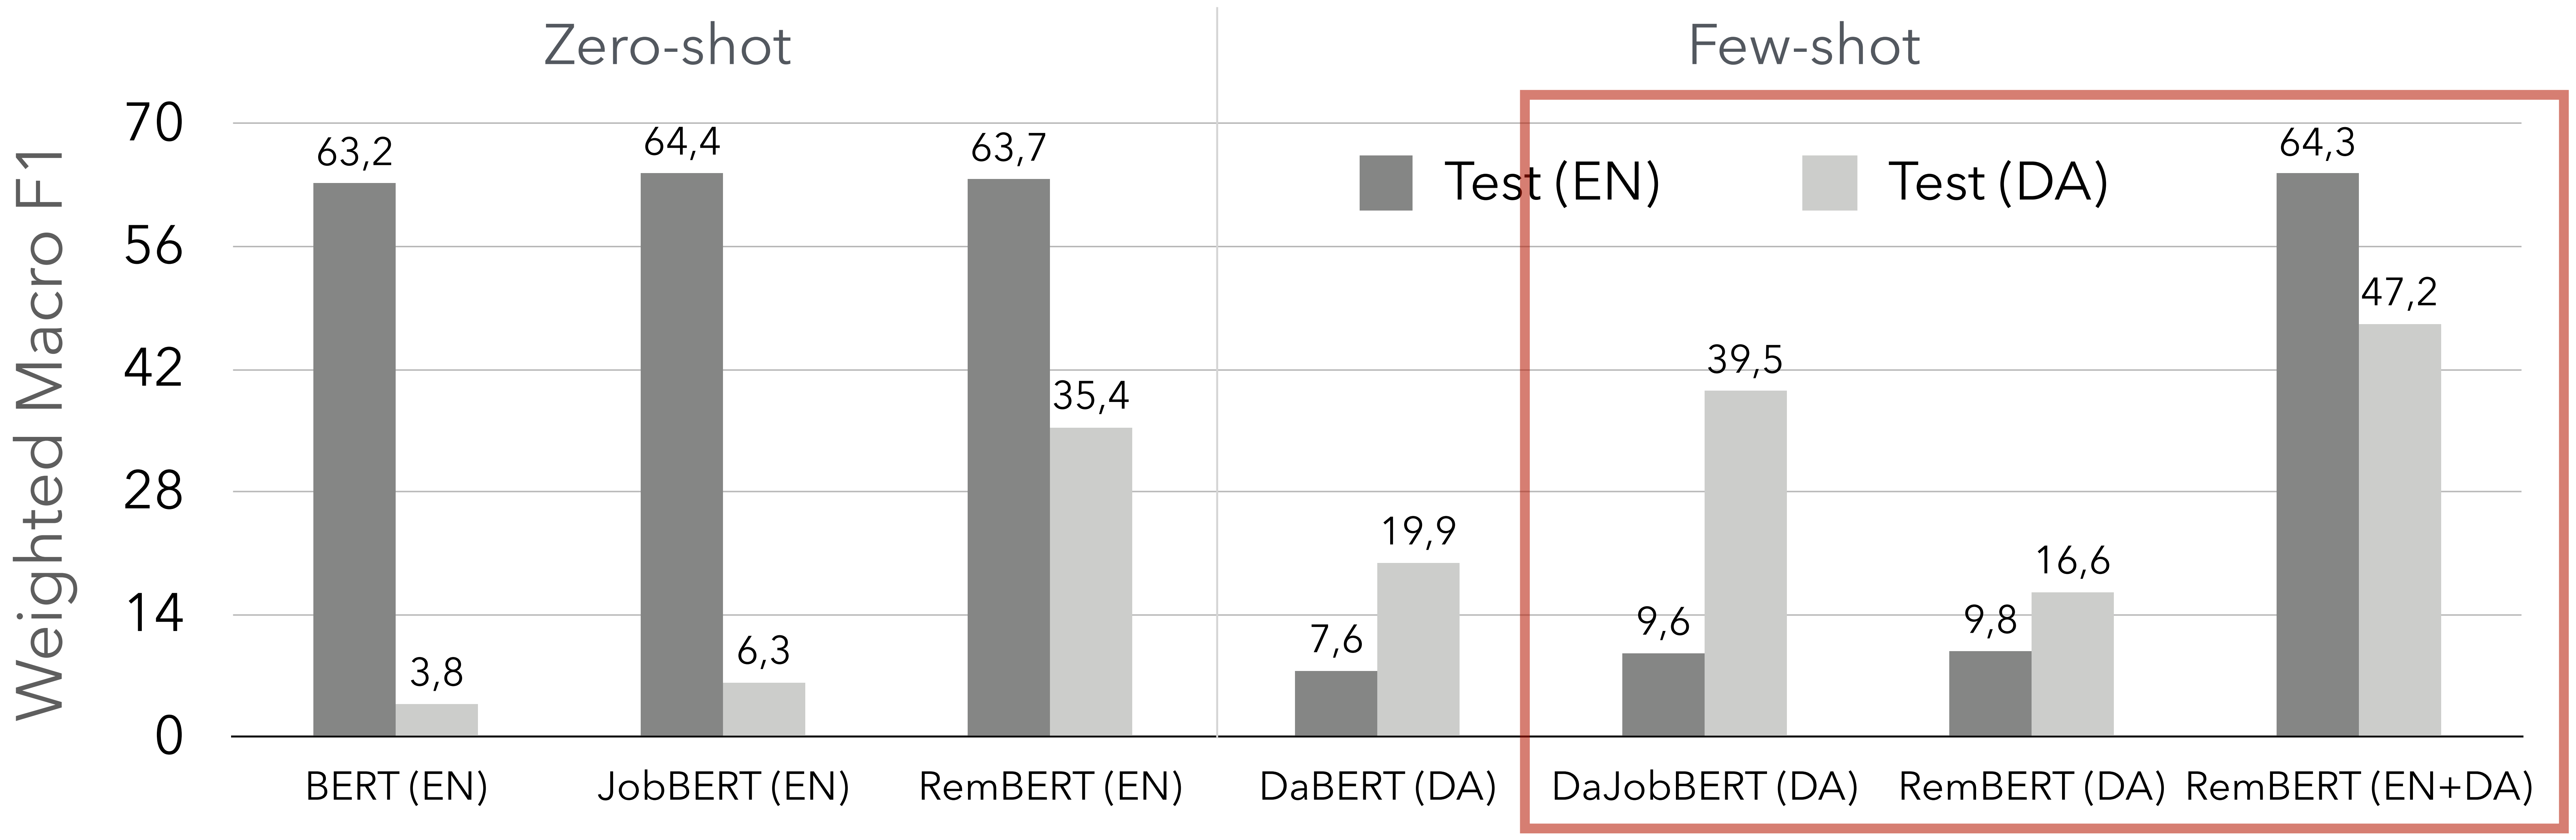

- Takeaway: Even with little in-language data, adding it to your larger pool of training data with other languages can significantly improve in-language downstream performance.

Takeaways

Kompetencer

- Distantly supervised ESCO skill labels are a cost-effective approach to train models.
- We again show that domain-adaptive pre-training helps, beyond English.
- Larger multilingual models are better than smaller specialized monolingual models.
- **Challenge 5, Output Space: If we use a taxonomy like ESCO as the labels, this results in a substantial output space during inference, adding a layer of complexity to the modeling process.**
    - All models perform well for classifying skill spans with 26 classes. >64 W. Macro-F1 for EN and >47 W. Macro-F1 for Danish.

Tackling the Long Tail

Tackling the Long Tail

Label distribution problem.

- In job descriptions, there is a *long-tail* pattern, popular skills are more commonly mentioned, while niche expertise appears infrequently.
- This results in a *sparsity of skills* in skill extraction datasets.
- E.g., one dataset contains a single annotation for "*identifying talent*" (S4.1.1.14), but other datasets do not contain this, can a model generalize?
  - Likely not.

NNOSE

NNOSE: Nearest Neighbor Occupational Skill Extraction

Mike **Zhang**, Rob **van der Goot**, Min-Yen **Kan**, and Barbara **Plank.**

To appear at EACL 2024.

A potential solution

NNOSE

- We explore Nearest Neighbor Language Models (NNLMs; Khandelwal et al., 2020), using the kNN algorithm as a retriever to retrieve context–token pairs from a datastore and leverage it during inference:
  - NNLMs have shown to memorize the training data better, aiding generalization;
  - NNLMs adapt to multiple domains without re-training;
  - NNLMs excels at predicting rare patterns, particularly in the long-tail.
- Concretely…

Overview of the idea

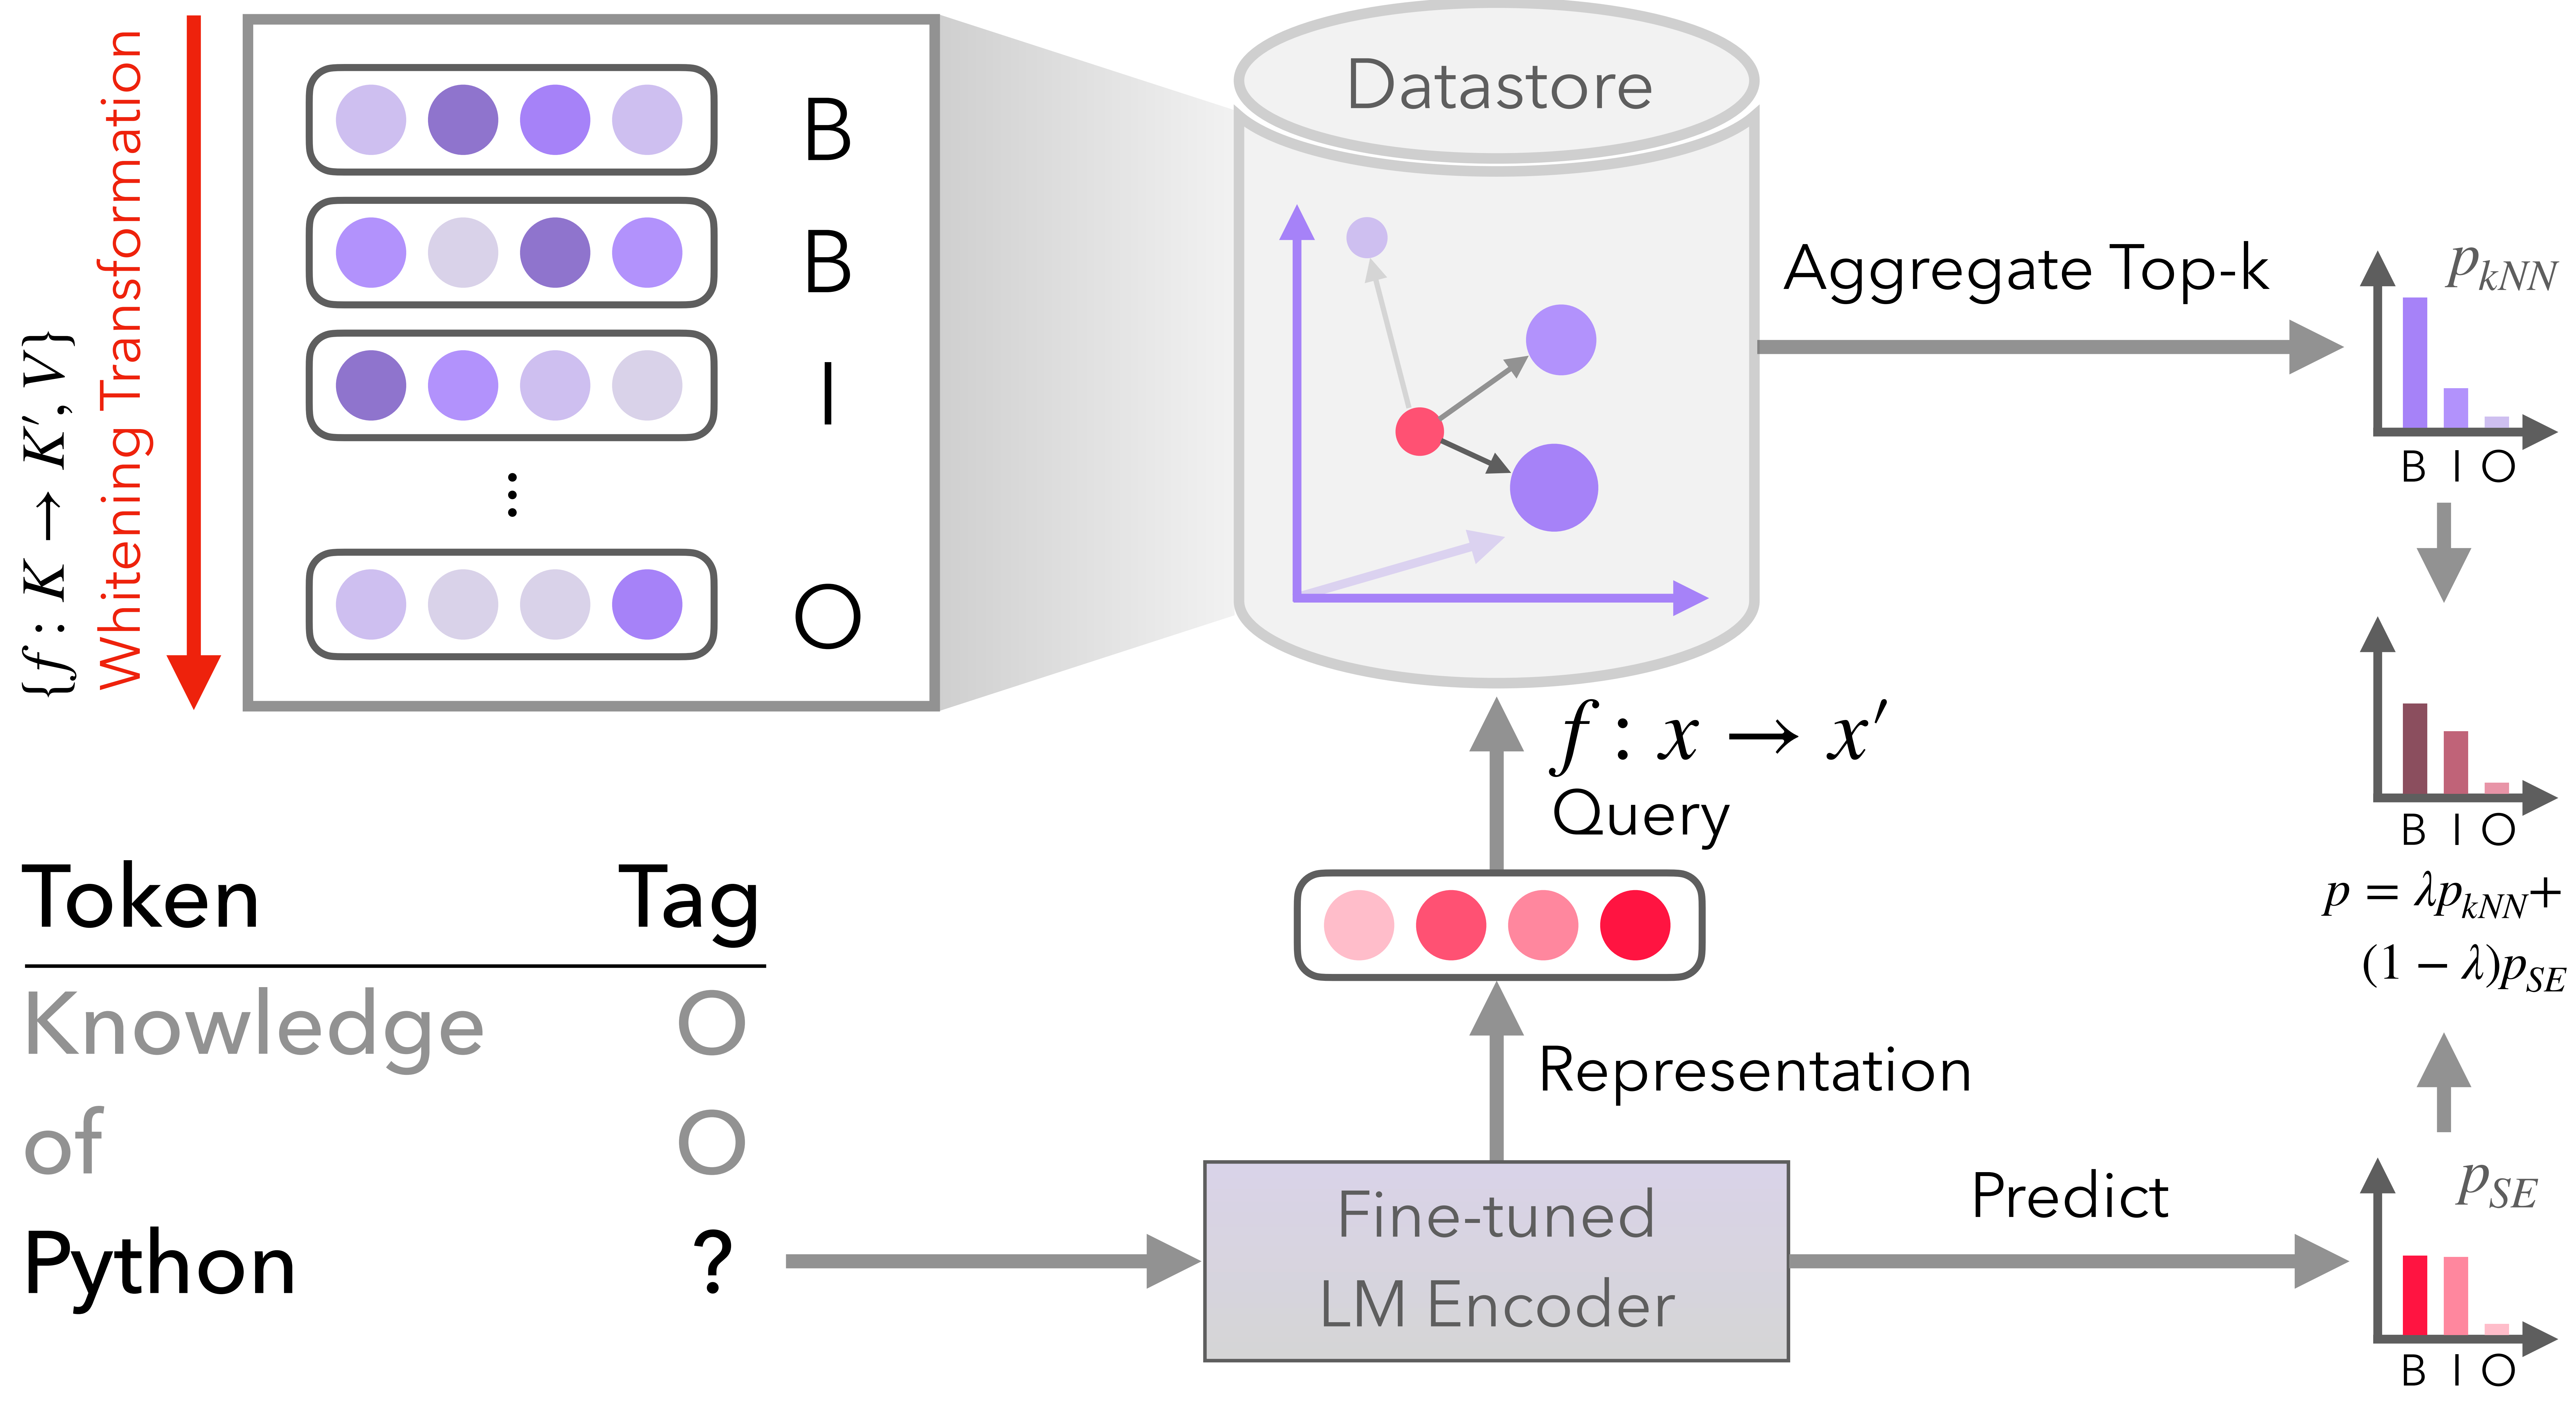

Setup

Models and Data

- Three models:
  - **JobBERT** (Zhang et al., 2022)
  - **RoBERTa** (Liu et al., 2019)
  - **JobBERTa** (This work), continuously pre-trained on 3.2M job posting sentences.
- Data:

Dataset	**Location**	**License**	**Train**	**Dev**	**Test**	**D (tokens)**
SkillSpan	*	CC-BY-4.0	5,866	3,992	4,680	86.5K
Sayfullina	UK	UNK	3,706	1,854	1,853	53.1K
Green	UK	CC-BY-4.0	8,670	963	336	209.5K
Total						348.2K

Table 1: **Dataset Statistics.** We provide statistics for all three datasets, including the location and license. Input granularity is at the token level, with performance measured in span-F1. The size of the datastore D is in tokens and determined by embedding tokens and their context from the **training** sets, resulting in approximately 350K keys.

Results

With and without kNN

- The best-performing baseline model is JobBERTa, achieving more than 4 points span-F1 improvement over JobBERT and 2 points higher than RoBERTa on average.

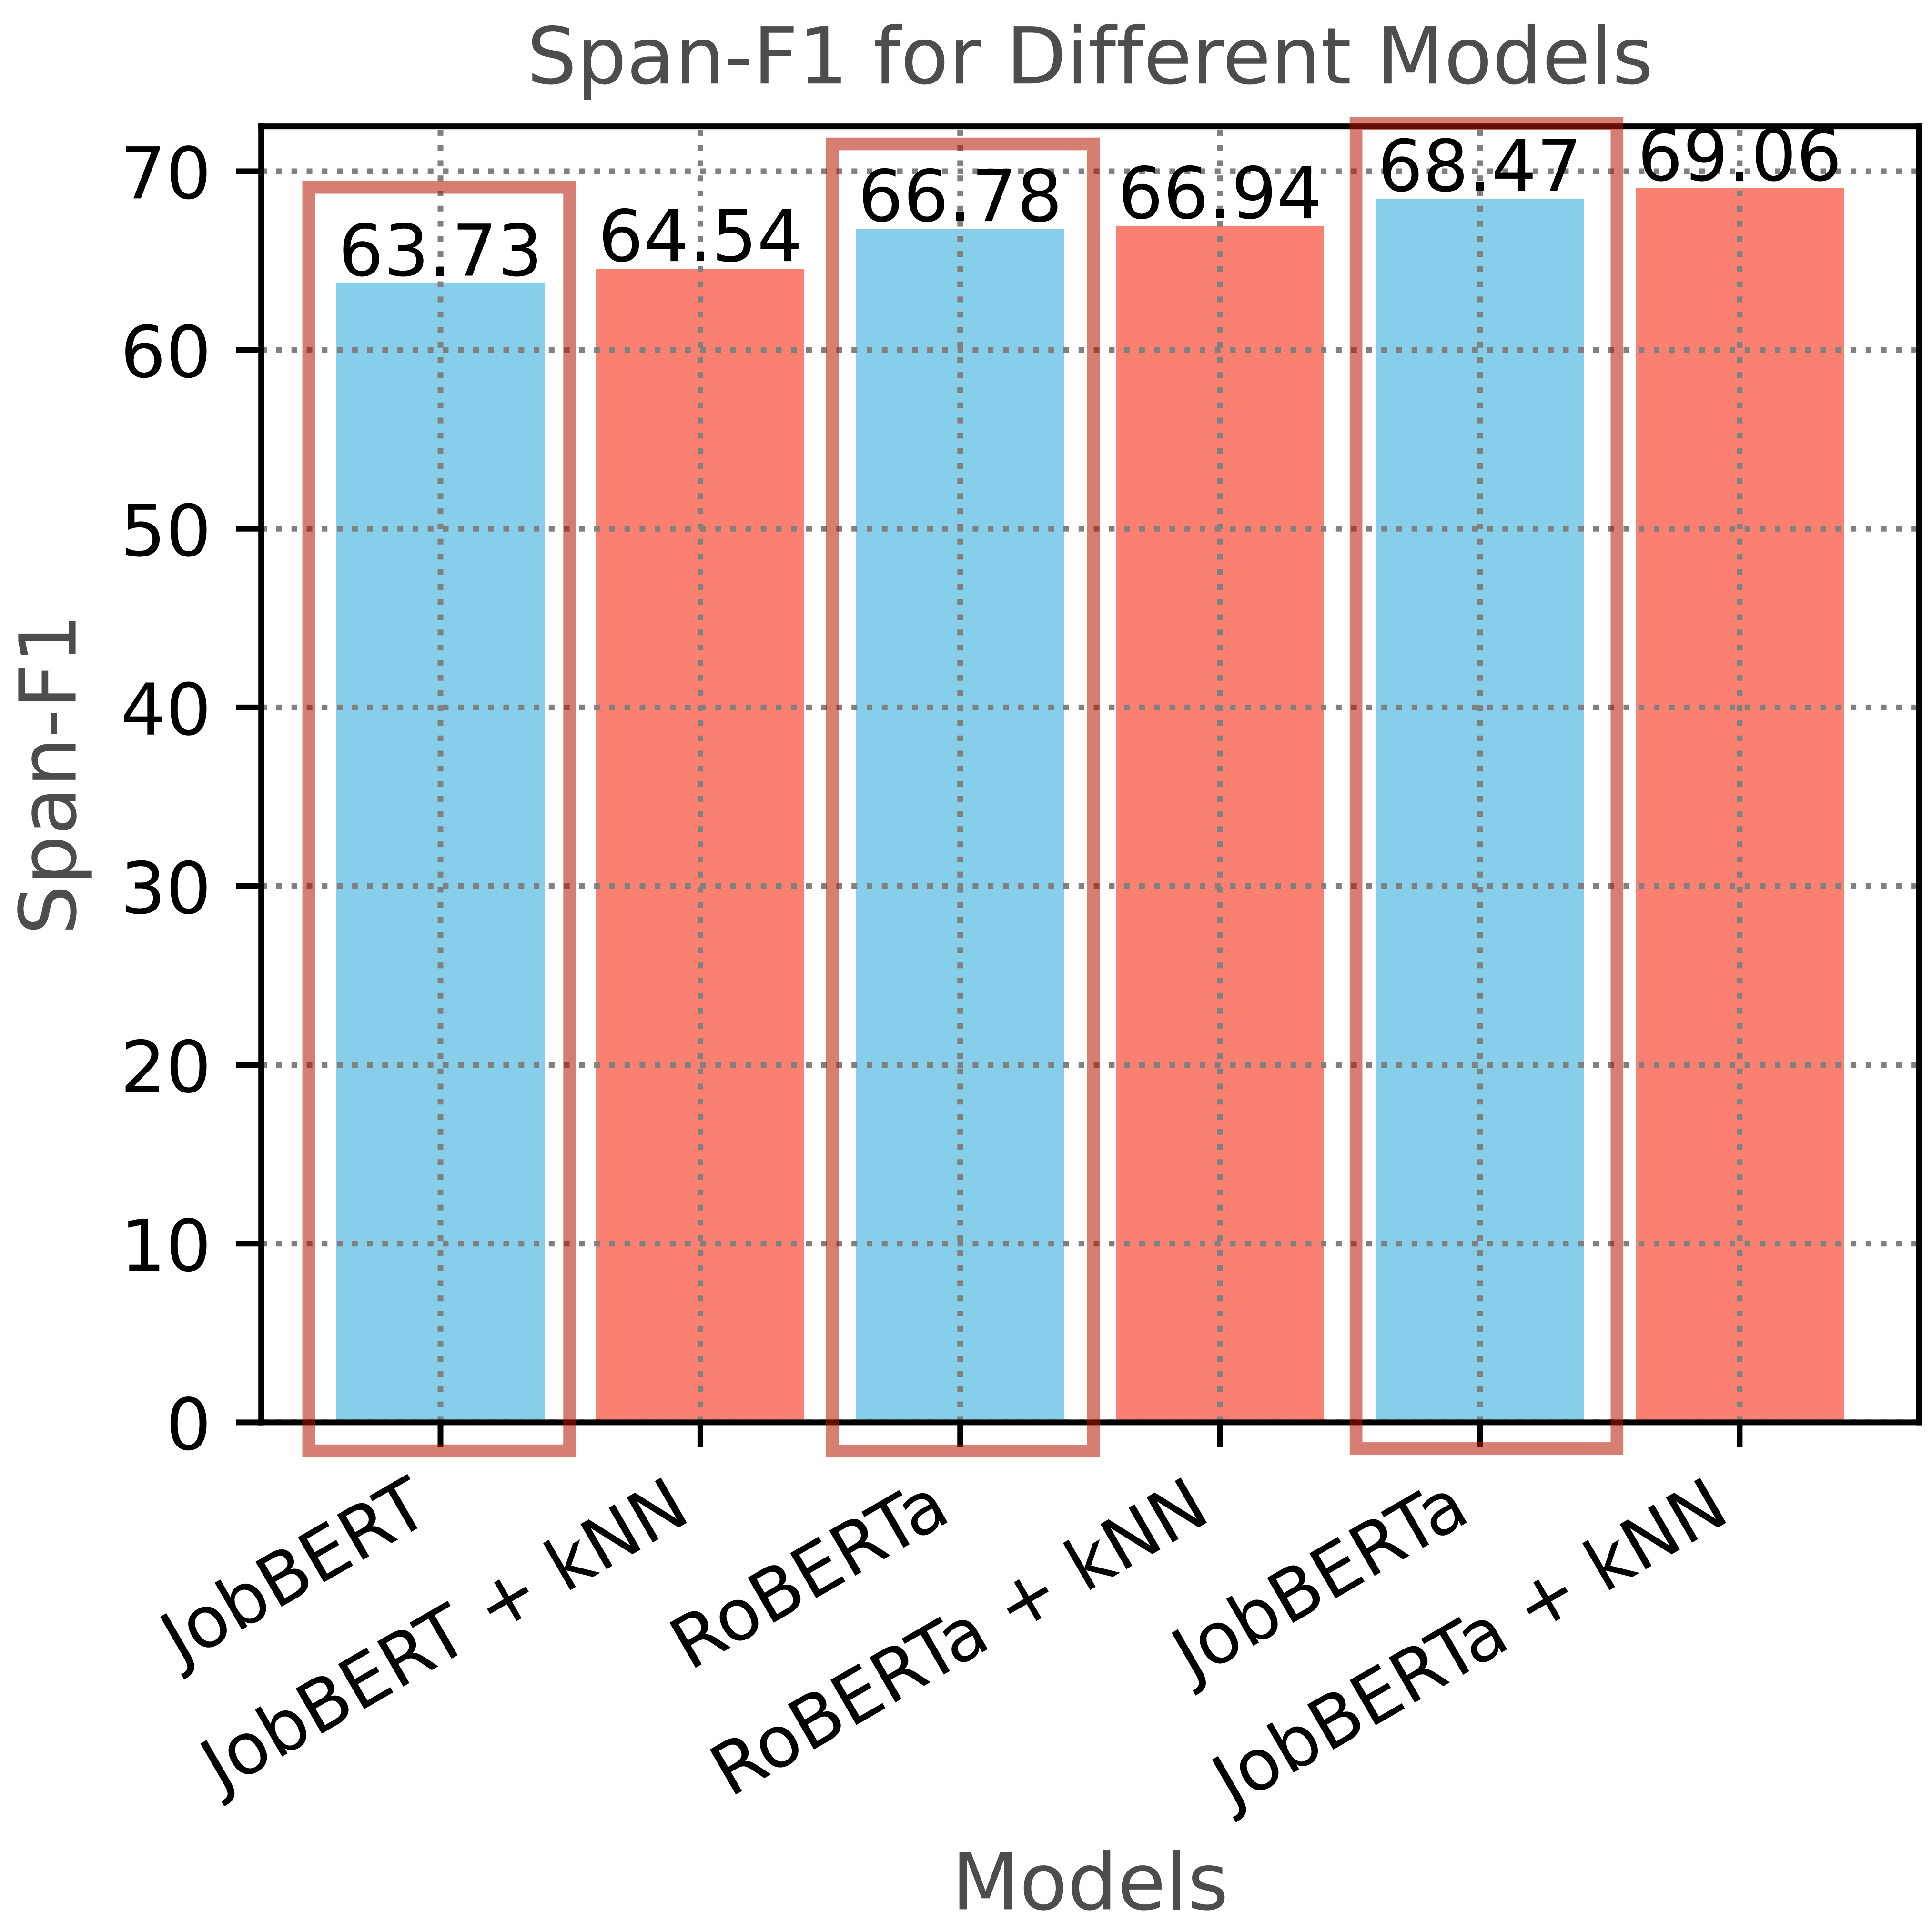

Results

With and without kNN

- The best-performing baseline model is JobBERTa, achieving more than 4 points span-F1 improvement over JobBERT and 2 points higher than RoBERTa on average.
- All models benefit from the NNOSE setup (+kNN), JobBERT and JobBERTa shows the largest improvements.

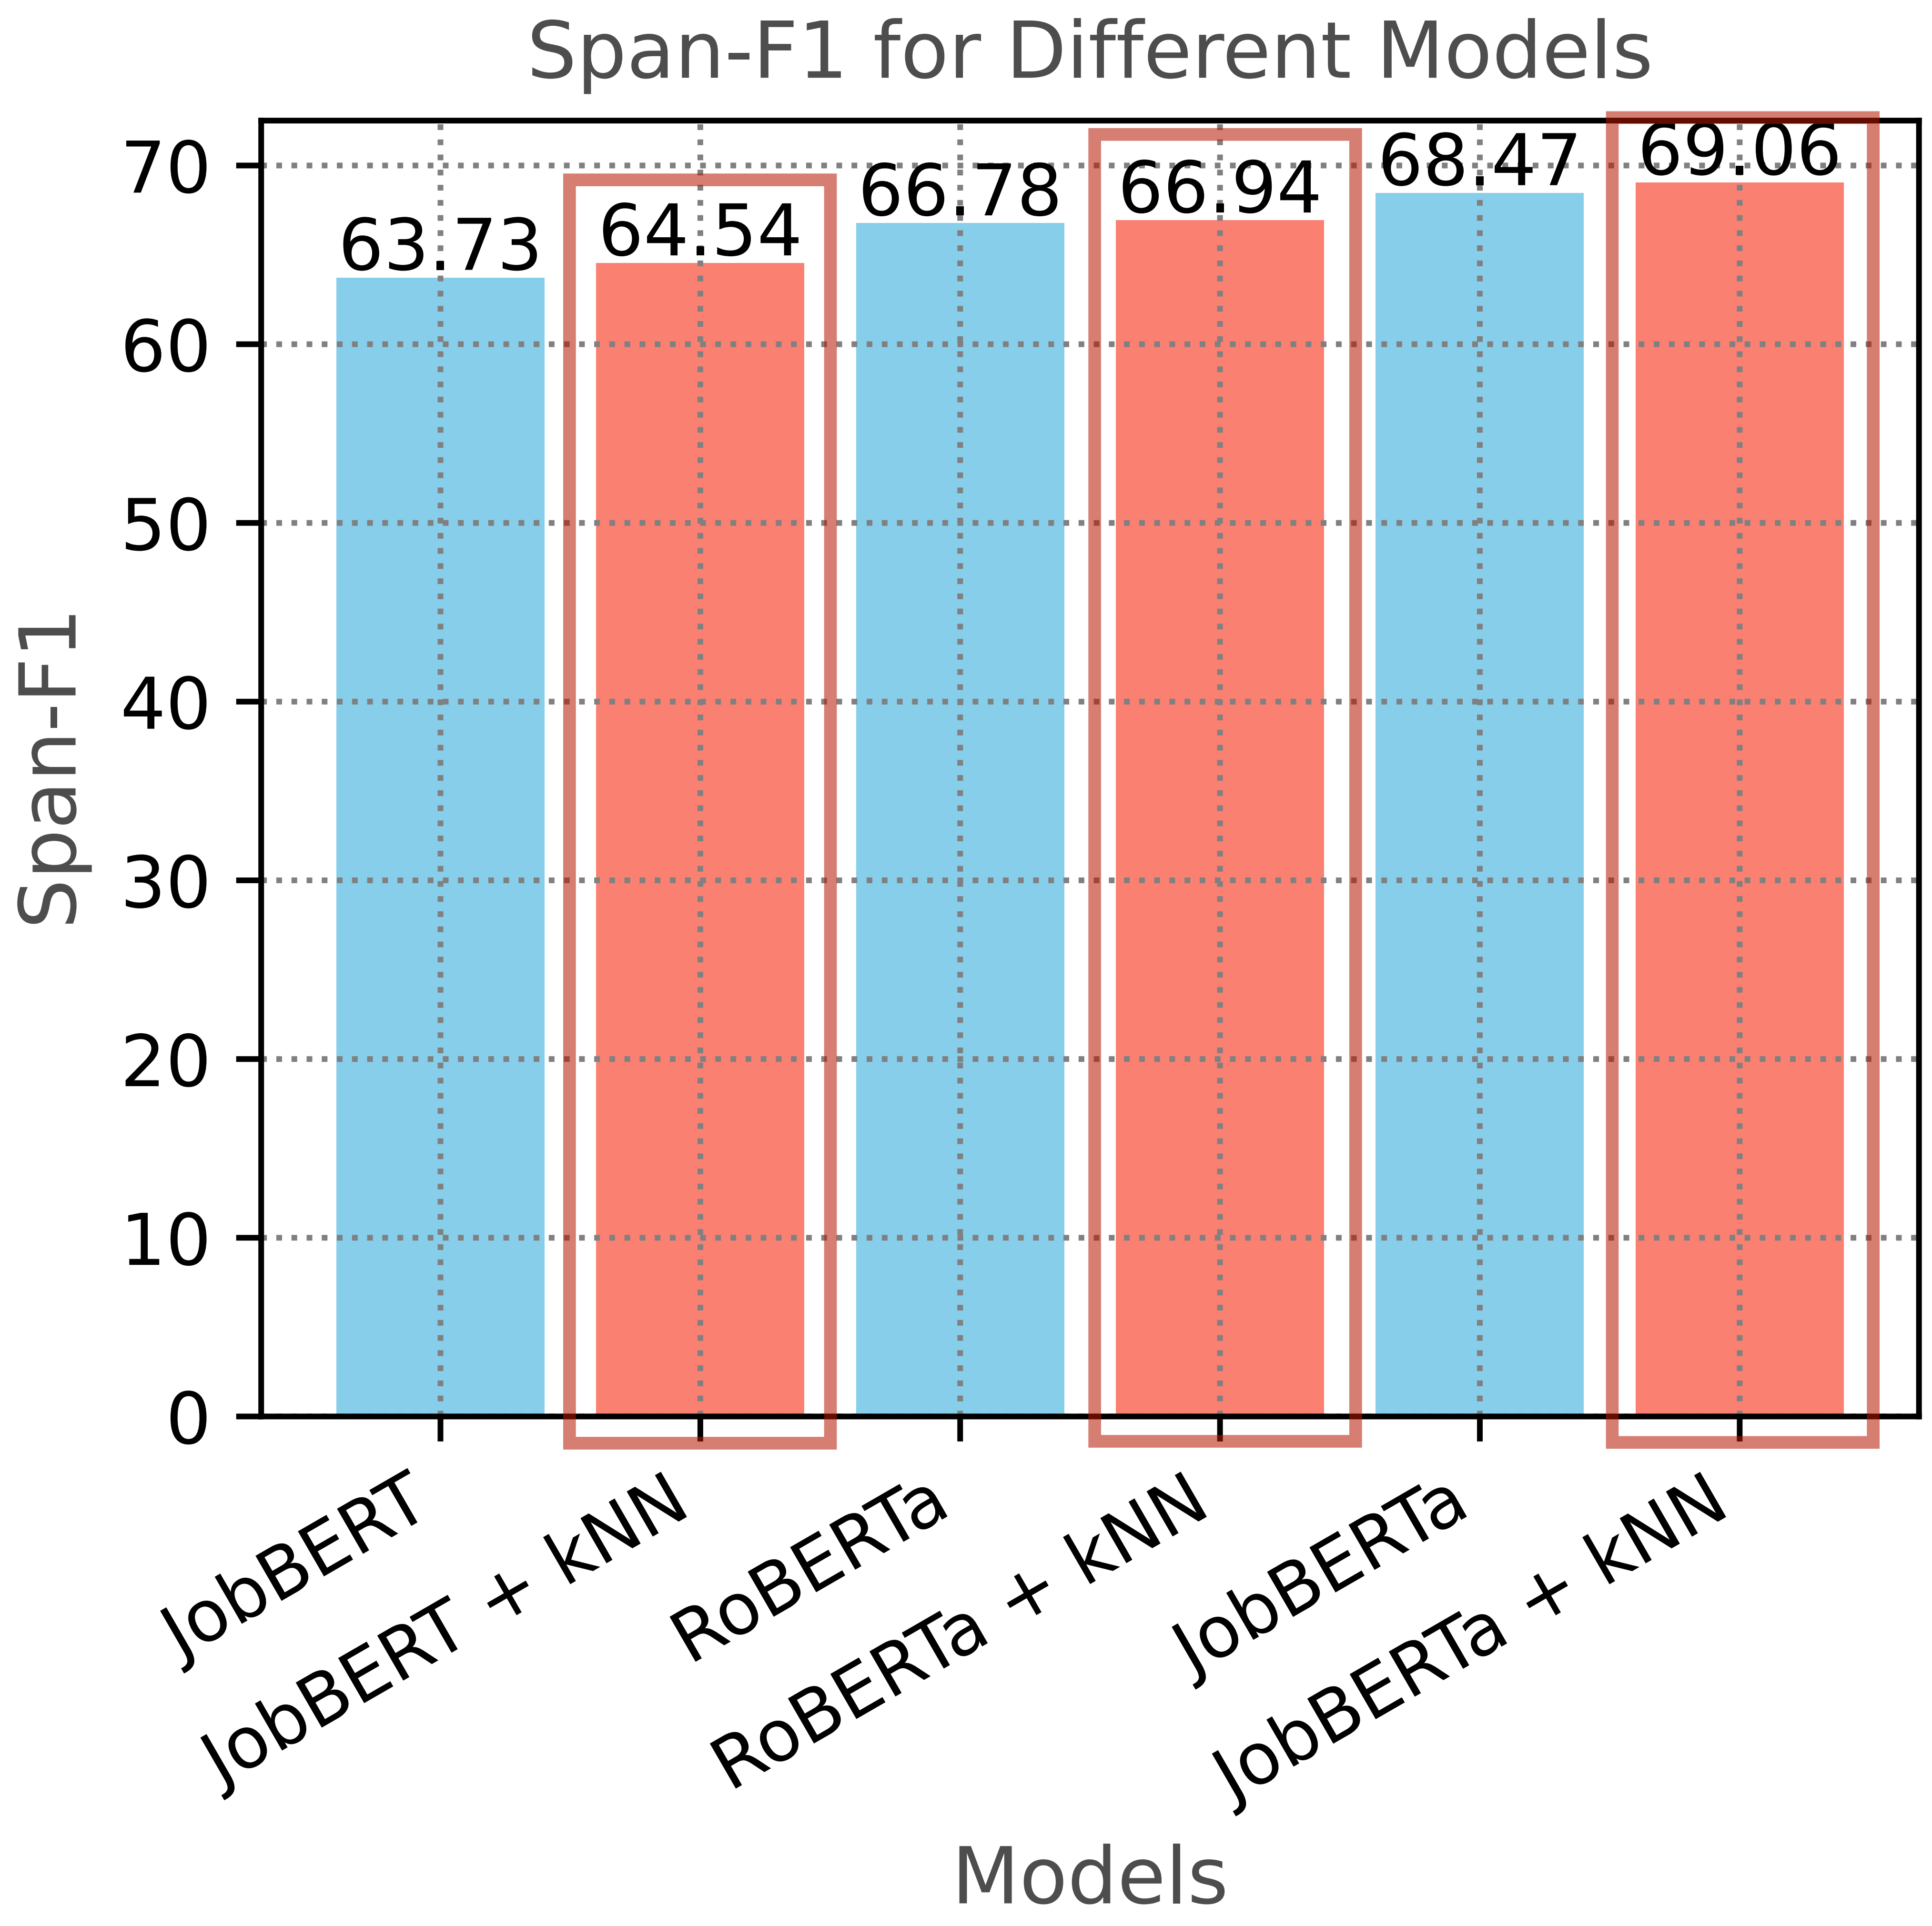

Analysis (1)

Cross-dataset Analysis

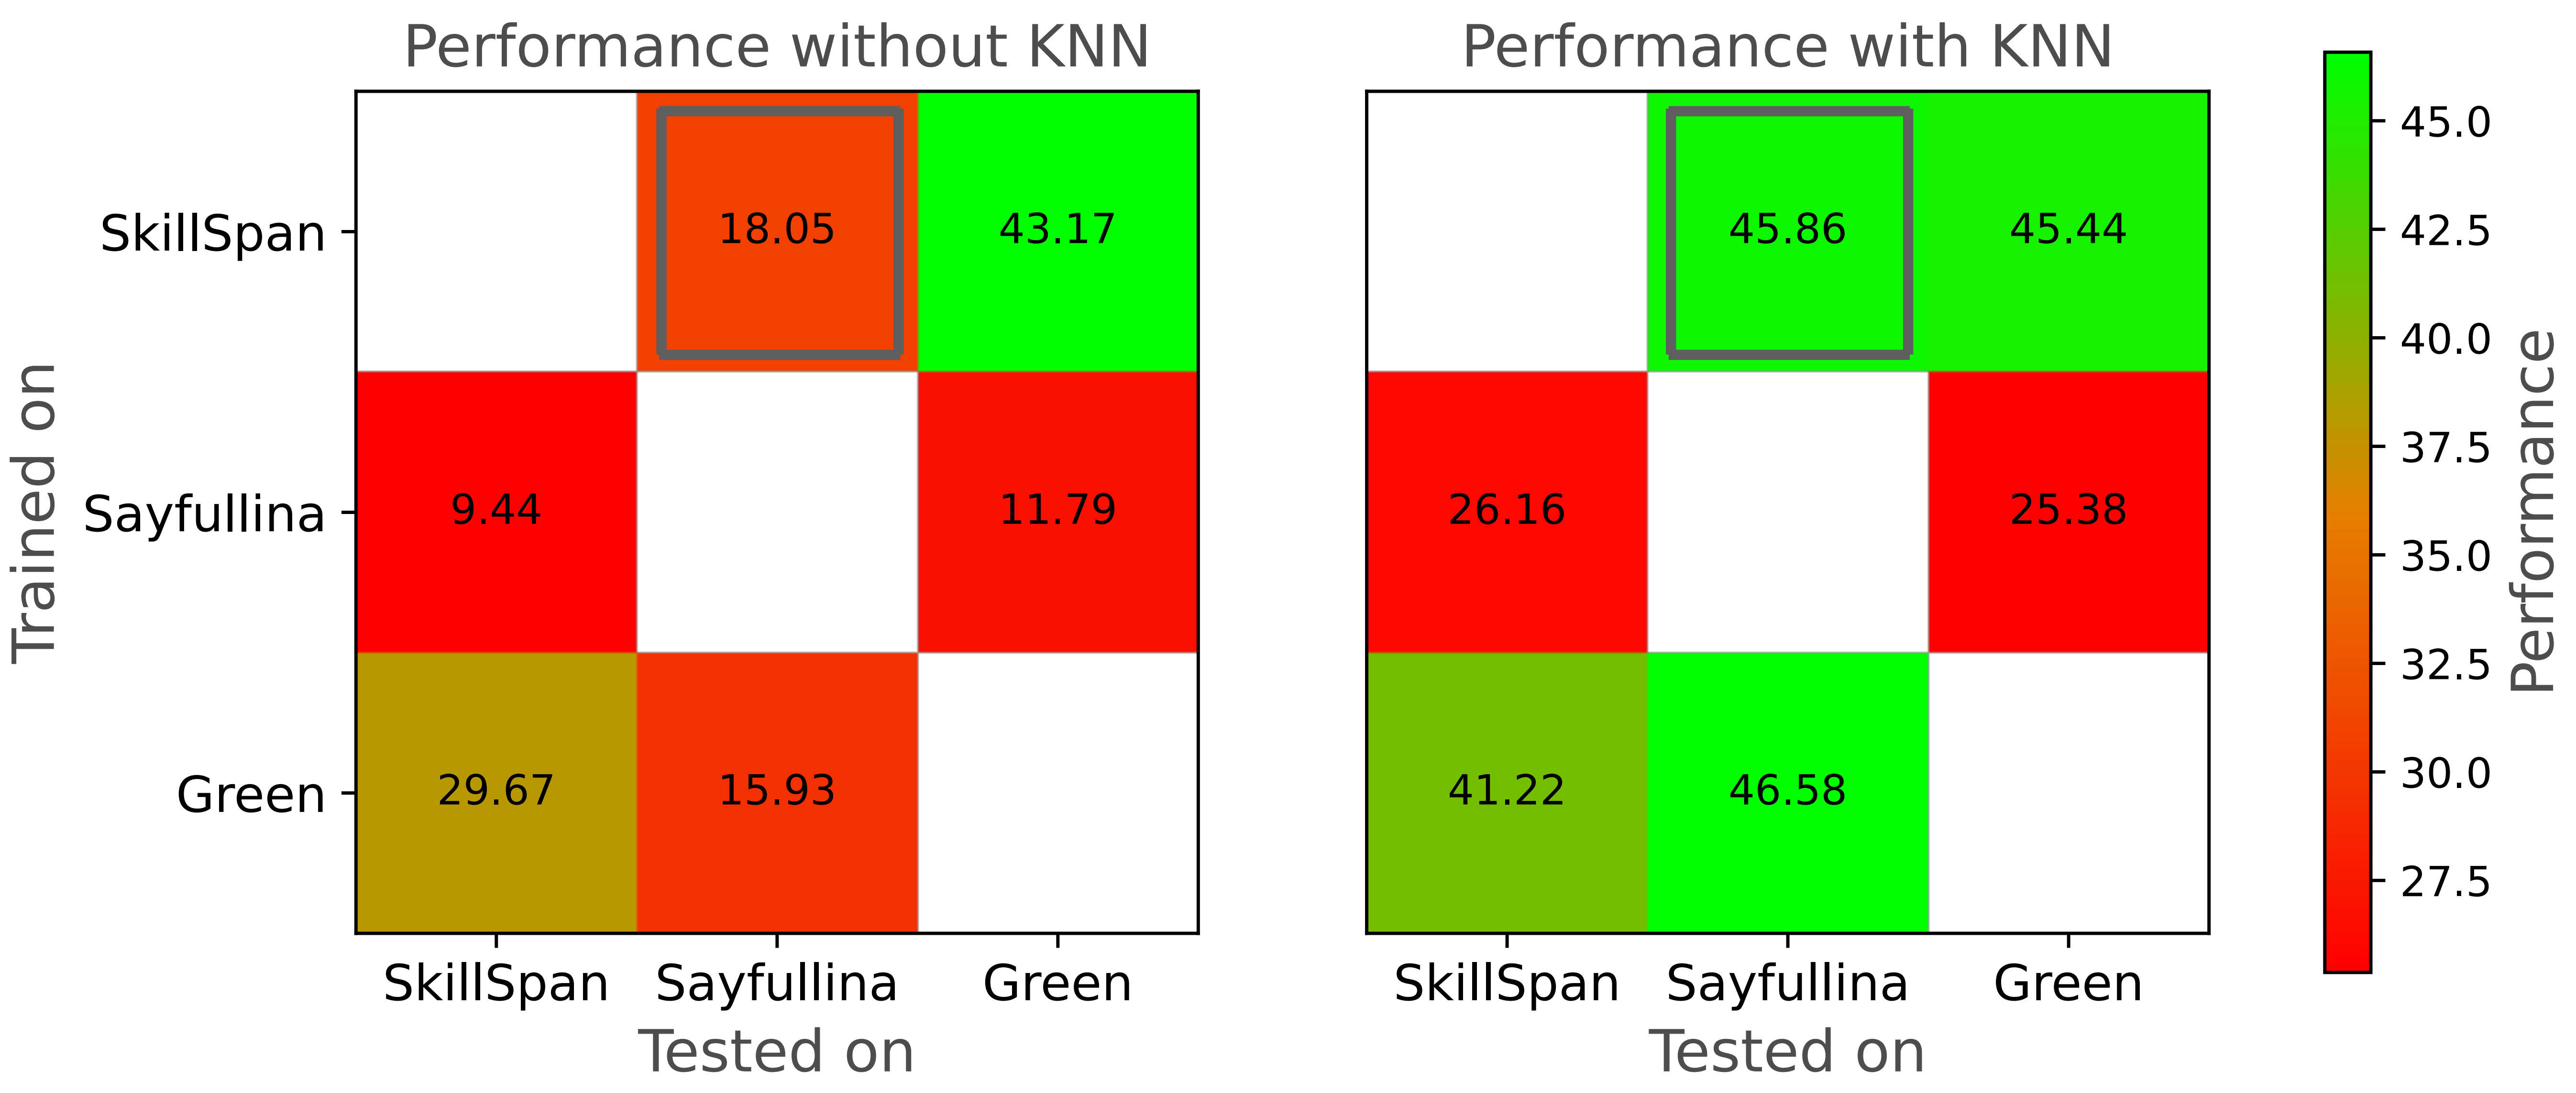

- We observe large gains in span-F1 using NNOSE in a cross-dataset setting.
  - Trained on dataset X, applied to dataset Y.
- We confirm findings similar to Khandelwal et al. (2020), that memorisation using NNLMs improves recall.

Analysis (2)

(In)frequent Patterns, in-dataset

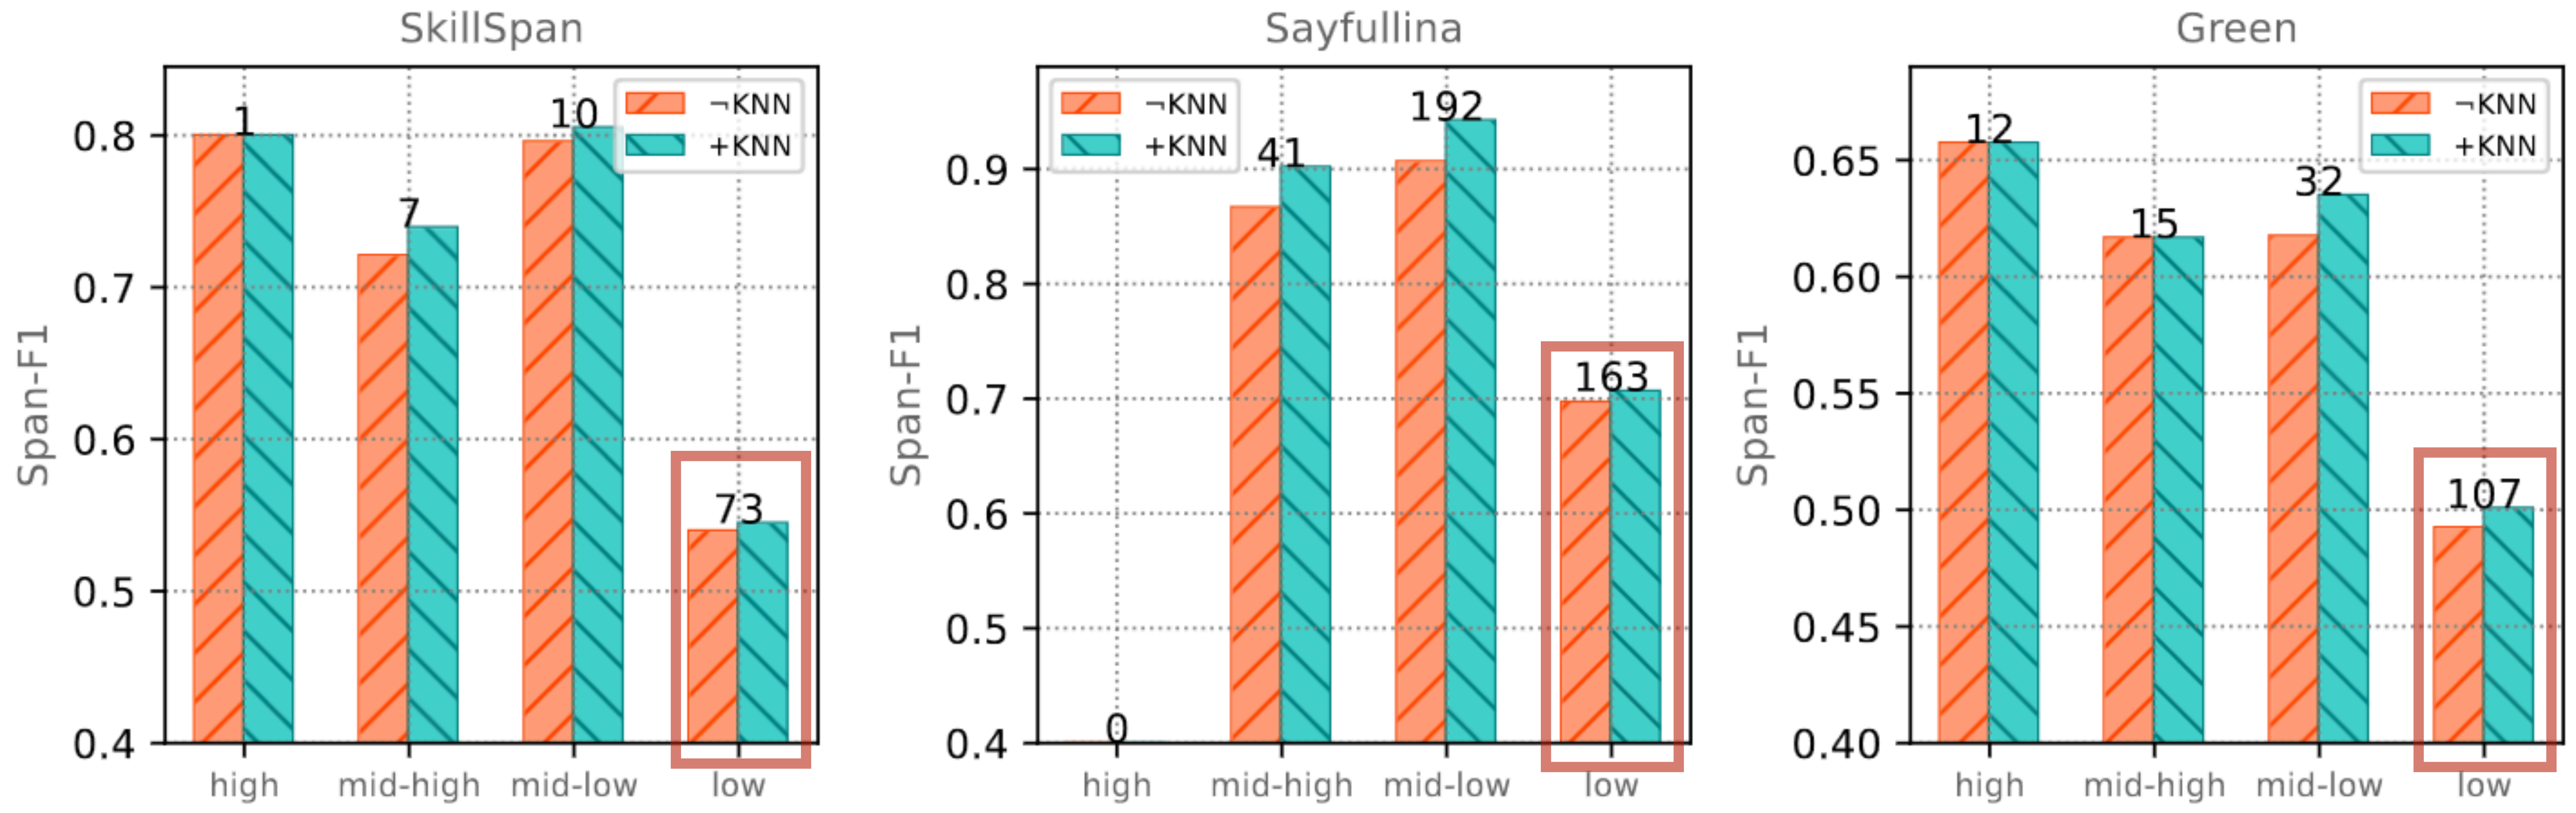

- Bin frequency of skills in the training set, ranging from 0–15 occurrences and is grouped into low, mid–low, mid–high, and high–frequency bins
  - (0–3, 4–6, 7–10, 10–15, respectively).
- Infrequent skills are the most difficult and make up the largest bucket, and our approach is able to improve on them on all three datasets.

Analysis (3)

(In)frequent Patterns, cross-dataset

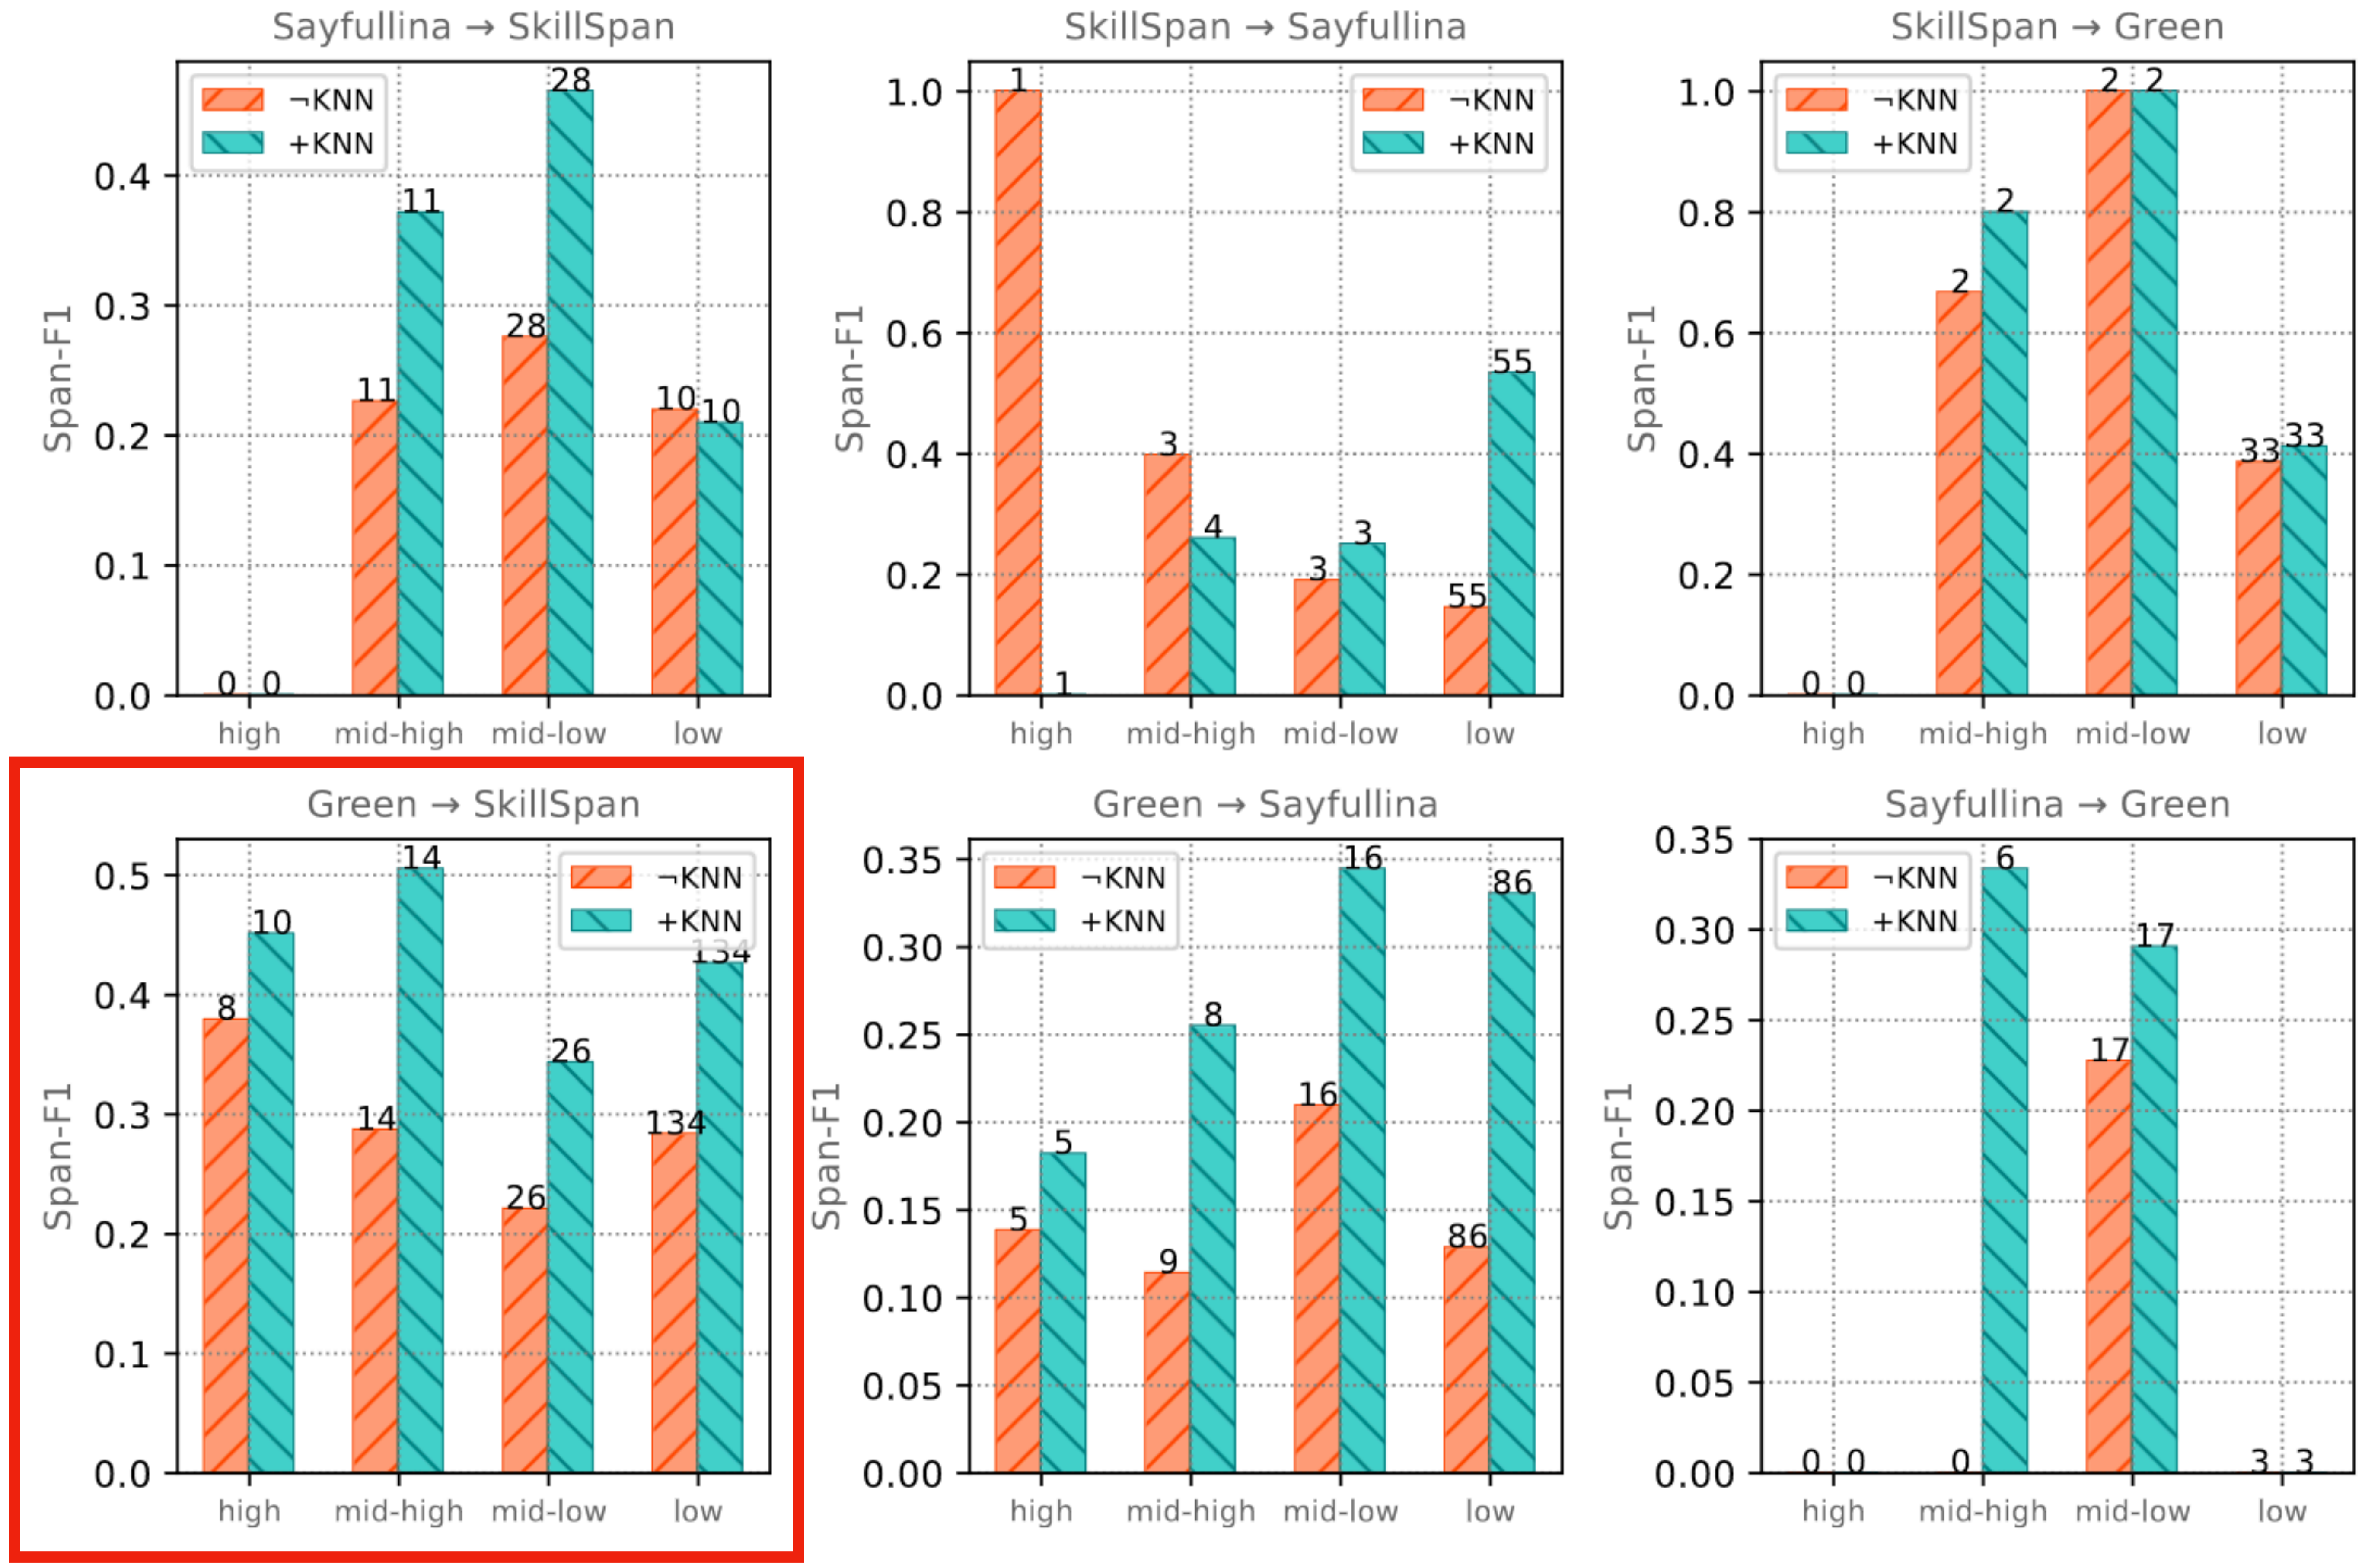

- Similar to the previous slide, we do the same analysis in the cross-dataset setting.
- In the cross-dataset setting, we observe a large gain in performance when using NNOSE.

Takeaways

NNOSE

- **Challenge 3, Long-tail: Certain skills are infrequently mentioned in a comprehensive set of job postings.**
    - NNOSE performs well on rare skills and also enhances the performance of more frequent skill patterns.
    - NNOSE improves on all fronts in a cross-dataset setting.
- We show that NNOSE is a promising approach for application-specific skill extraction setups (e.g., Retrieval Augmentation).

Linking Skills to Existing Resources

Linking Skills to Existing Resources

How can we make the extraced skills useful?

- Previously, we have shown to be able to classify spans into their ESCO taxonomy code (26 classes).
  - Not taking surrounding context into account

- How can we scale this up to ~13,000 classes?

EL-Job

Entity Linking in the Job Market Domain

Mike **Zhang**, Rob **van der Goot**, and Barbara **Plank.**

To appear at Findings of EACL 2024.

Setting Context

Skill Matching == Entity Linking (Wu et al., 2020; Discriminative)

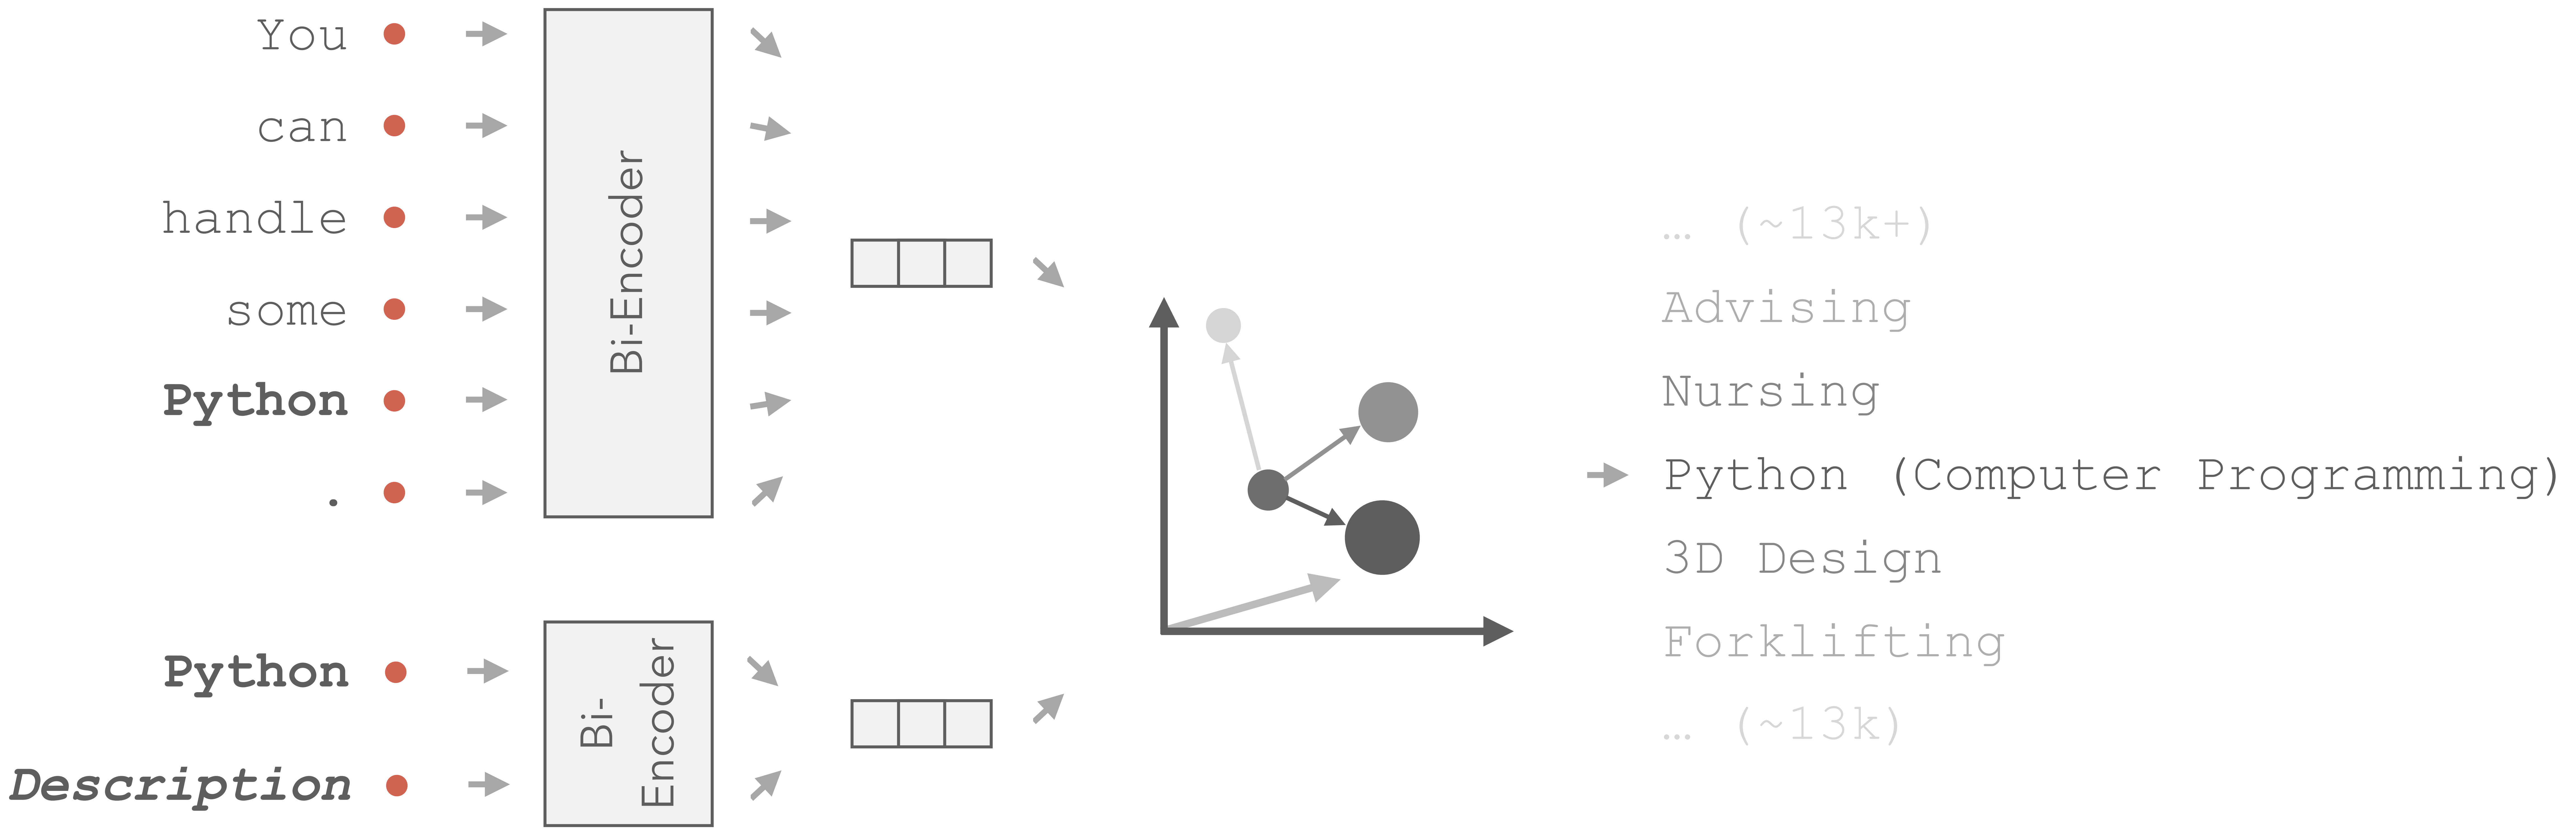

Setting Context

Generative (de Cao et al. 2021)

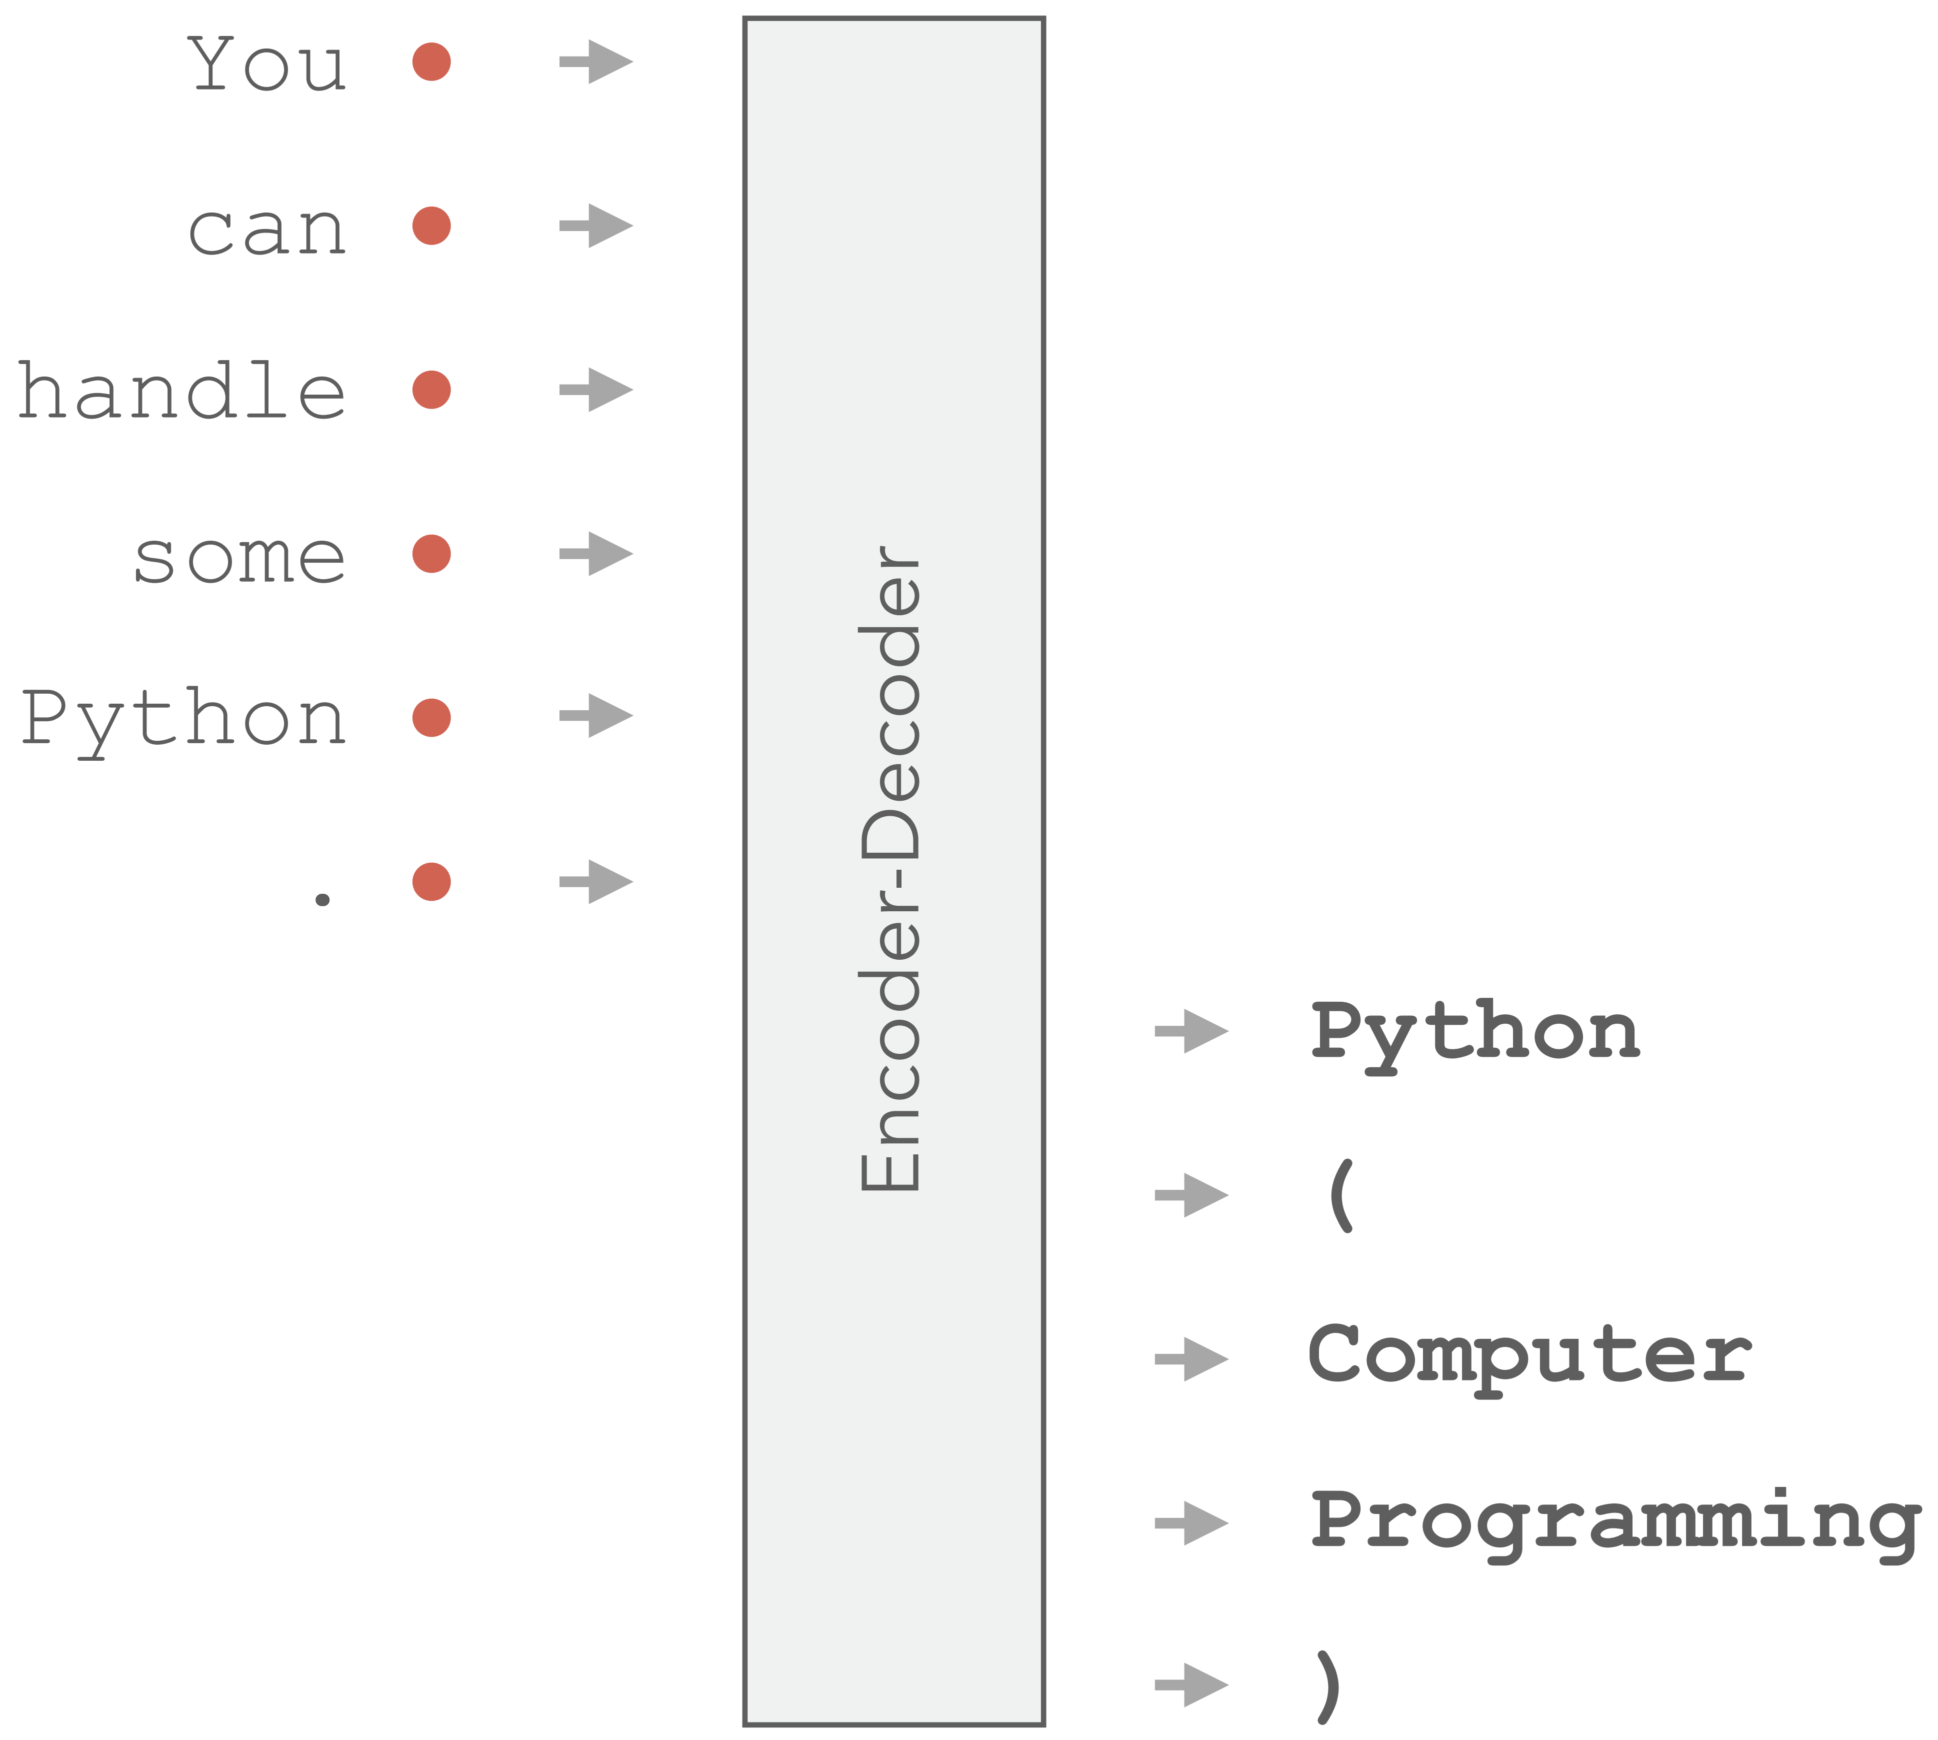

Setting Context

Generative, GIF from de Cao (2021).

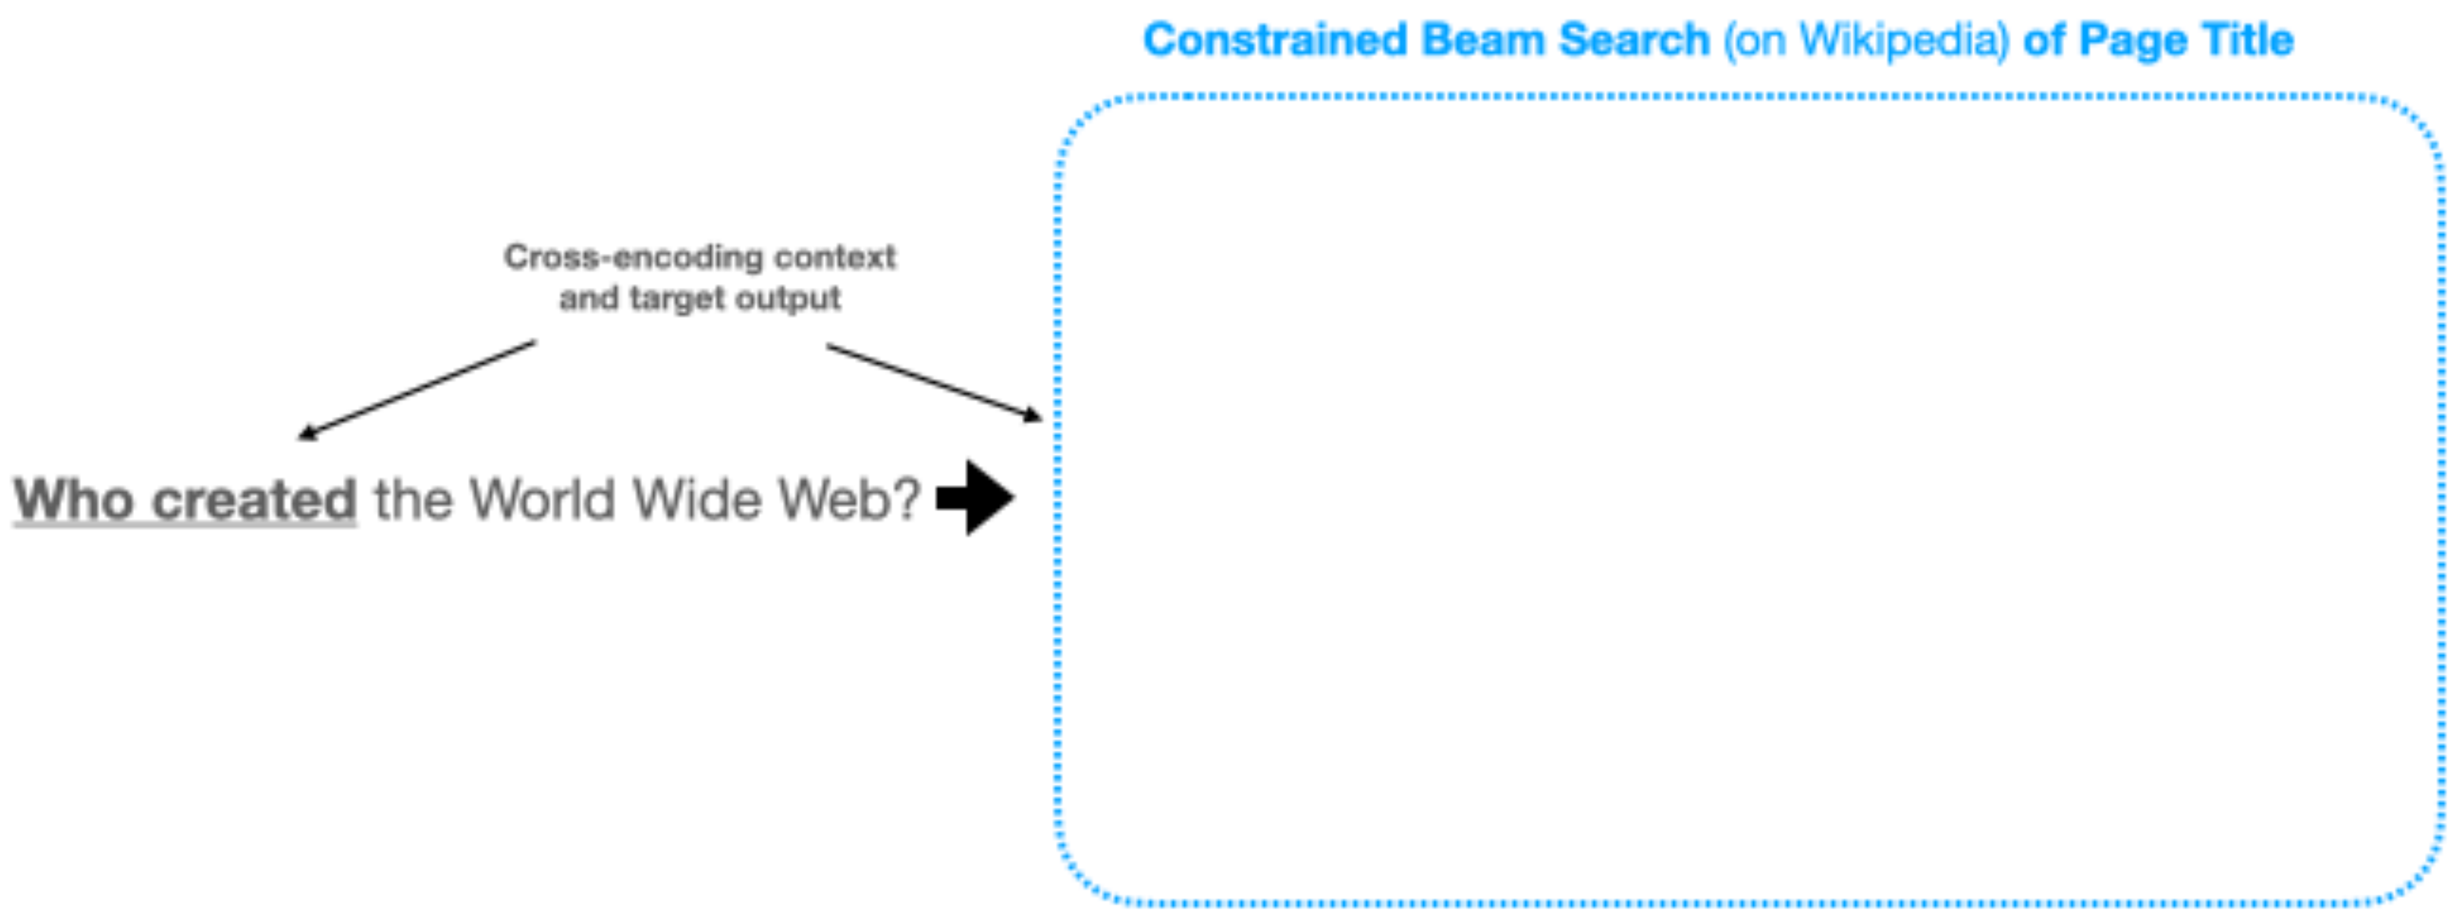

Setup

Two models and a benchmark

- Training data from Decorte et al. (2023), a synthetic training set generated using GPT-4.
- Dev. and Test are human-curated, Decorte et al. (2022) built on top of Zhang et al. (2022) by manually matching skill spans to the ESCO taxonomy.
- Models, two entity-linkers:
  - BLINK (Wu et al., 2020), discriminative
  - GENRE (Cao et al., 2021), generative

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

Results of Linking

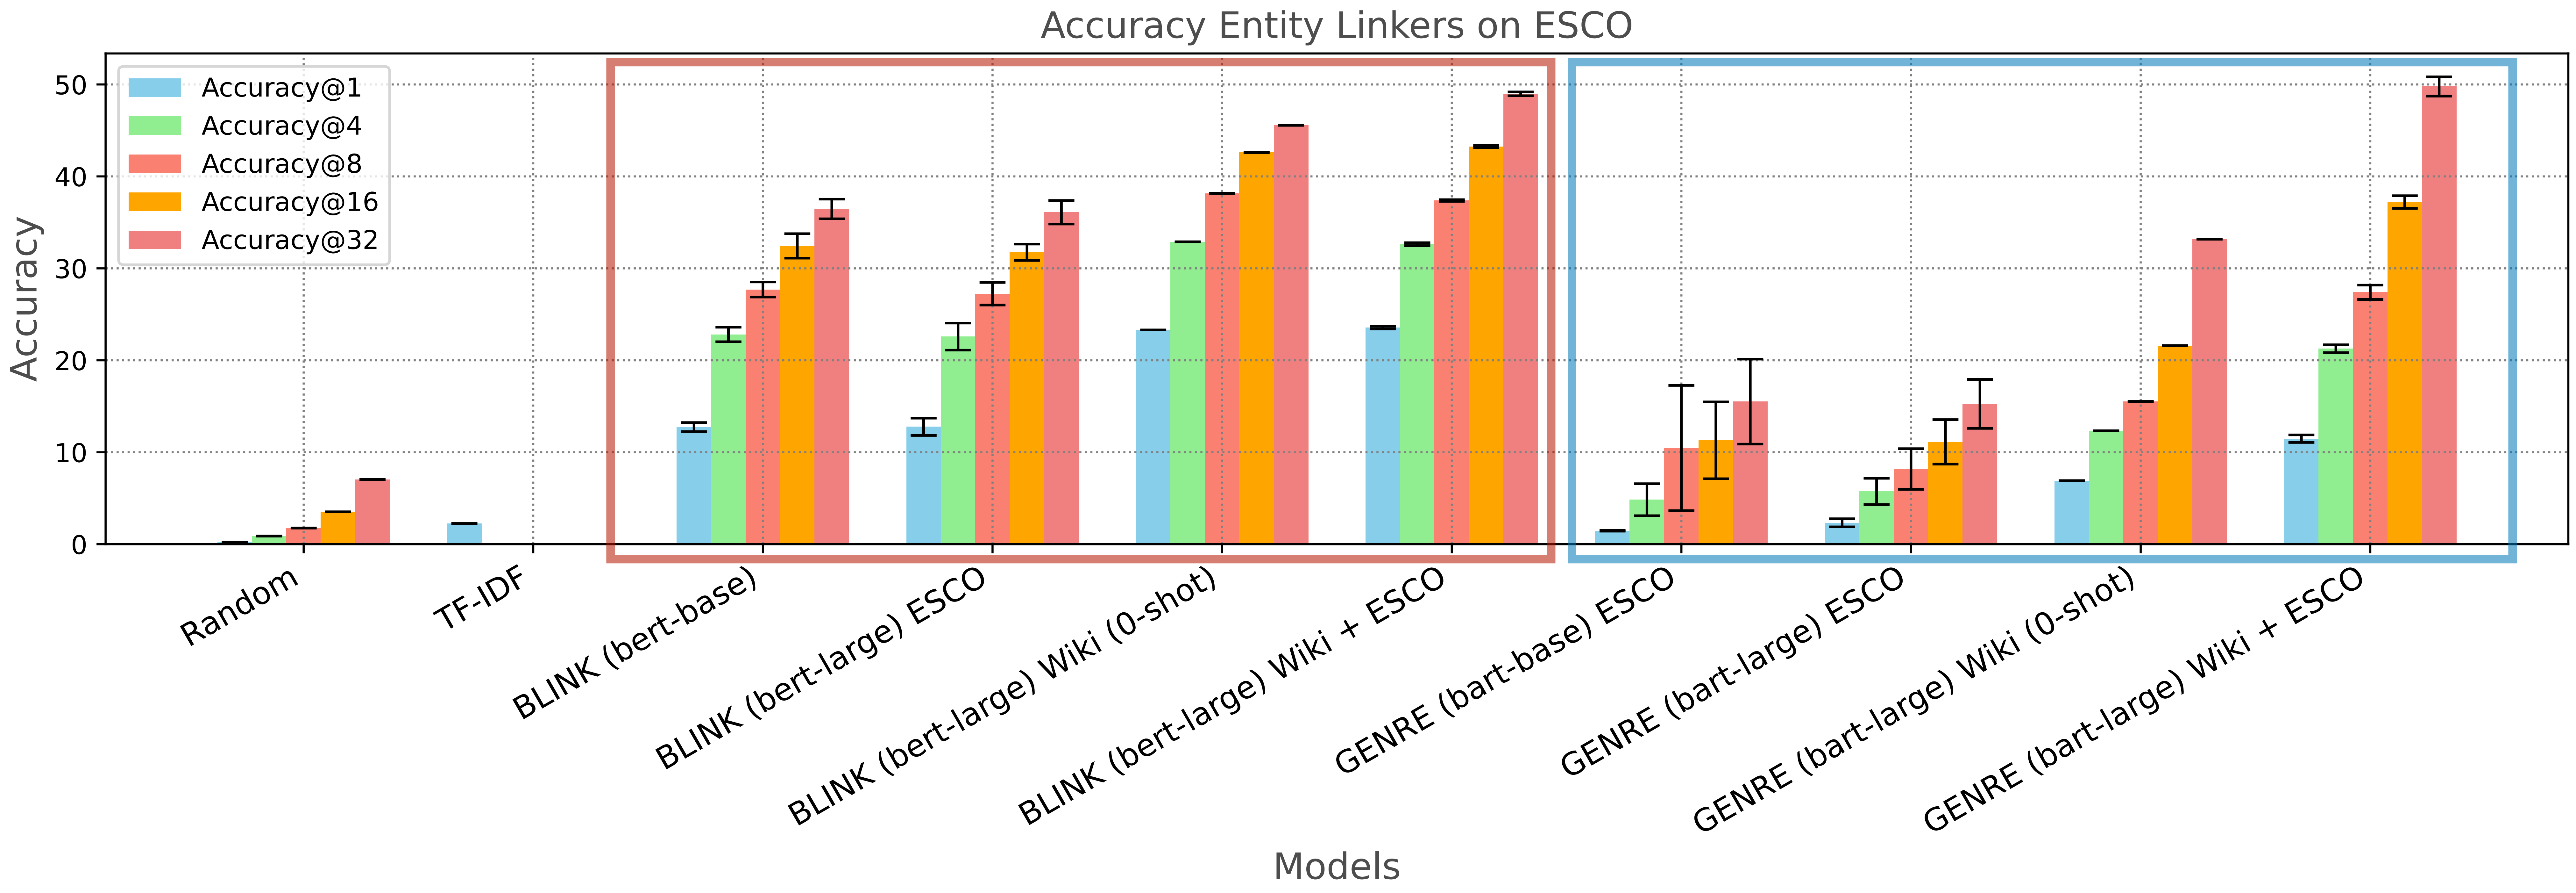

- We show that the discriminative linker (i.e., BLINK) is more suited for this task.
- Strict Linking evaluation favours BLINK, however if we relax the number of candidates (i.e., higher # of k), we observe that GENRE performs better.

Qualitative Analysis

Challenge 1: What does it link?

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

Takeaways

EL in Job Market

- We show that the discriminative linker (i.e., BLINK) is more suited for this task.
- Strict Linking evaluation favours BLINK, however if we relax the number of candidates (i.e., higher # of k), we observe that GENRE performs better.
- Qualitative analysis shows that both models link valid predictions for a skill mention.
  - However, entity linking does not allow for multiple correct predictions, highlighting the need for a more comprehensive evaluation.

- **Challenge 1: Entity Linking models have shown to be able to link <u>implicit skills</u>.**
- **Challenge 5: There is a reasonable performance linking skill spans to ~13k labels.**

Conclusions

Conclusions

What can we learn from this?

- Skill Extraction is a hard task with many interesting challenges that could be of interest for general NLP research:
  - Long spans; Discontinuity; Implicitness; Long-tail extraction; Large label space.
- We show four contributions tackling four challenges in Skill Extraction.
  - CH1: Entity Linking models have shown to be able to link implicit skills;
  - CH2: Language models are able to predict long skills on a reasonable level (up to 7 tokens);
  - CH3: Retrieval-augmentation helps predicting skills in the long tail;
  - CH5: Models are able to match skills to their taxonomy counterpart, but need re-ranking methods to perform better.

Future Directions

Where do we go?

- **Beyond Job Postings**
  - Other type of document (e.g., resumes).
- **Large-scale End-to-End datasets**
  - Multiple span-level annotations of skills and linking them directly to a taxonomy code.
  - Synthetic data?
- **Large Language Models**
  - Conjoined skills can likely be split up by LLMs.
  - Constraining output space.


Thanks for listening!
Questions?
Max™

References

Rauner, F., & Maclean, R. (Eds.). (2008). Handbook of technical and vocational education and training research (Vol. 49). Dordrecht: Springer.

Zhao, Z., & Rauner, F. (Eds.). (2014). Areas of vocational education research. Dordrecht, Heidelberg: Springer.

le Vrang, M., Papantoniou, A., Pauwels, E., Fannes, P., Vandensteen, D., & De Smedt, J. (2014). Esco: Boosting job matching in europe with semantic interoperability. Computer, 47(10), 57-64.

Zhang, M., Jensen, K., Sonniks, S., & Plank, B. (2022, July). SkillSpan: Hard and Soft Skill Extraction from English Job Postings. In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 4962-4984).

Caruana, R. (1997). Multitask learning. *Machine learning*, *28*, 41-75.

Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020, July). Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 8342-8360).

Mintz, M., Bills, S., Snow, R., & Jurafsky, D. (2009, August). 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 (pp. 1003-1011).

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.

Mike Zhang, Kristian Nørgaard Jensen, and Barbara Plank. 2022. Kompetencer: Fine-grained Skill Classification in Danish Job Postings via Distant Supervision and Transfer Learning. In Proceedings of the Thirteenth Language Resources and Evaluation Conference, pages 436–447, Marseille, France. European Language Resources Association.

Chung, H. W., Fevry, T., Tsai, H., Johnson, M., & Ruder, S. (2020, October). Rethinking Embedding Coupling in Pre-trained Language Models. In International Conference on Learning Representations.Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. arXiv preprint arXiv:2303.10130.

Su, J., Cao, J., Liu, W., & Ou, Y. (2021). Whitening sentence representations for better semantics and faster retrieval. arXiv preprint arXiv:2103.15316.

Yin, W., & Shang, L. (2022, December). Efficient Nearest Neighbor Emotion Classification with BERT-whitening. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (pp. 4738-4745).

Khandelwal, U., Levy, O., Jurafsky, D., Zettlemoyer, L., & Lewis, M. (2019, September). Generalization through Memorization: Nearest Neighbor Language Models. In *International Conference on Learning Representations*.

Wu, L., Petroni, F., Josifoski, M., Riedel, S., & Zettlemoyer, L. (2020, November). Scalable Zero-shot Entity Linking with Dense Entity Retrieval. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 6397-6407).

De Cao, N., Izacard, G., Riedel, S., & Petroni, F. (2020, October). Autoregressive Entity Retrieval. In *International Conference on Learning Representations*.