

Efficient Domain Adaptation of Sentence Embeddings Using Adapters

Tim Schopf, Dennis N. Schneider, and Florian Matthes

Technical University of Munich, Department of Computer Science, Germany

{tim.schopf, dennis.schneider, matthes}@tum.de

Abstract

Sentence embeddings enable us to capture the semantic similarity of short texts. Most sentence embedding models are trained for general semantic textual similarity tasks. Therefore, to use sentence embeddings in a particular domain, the model must be adapted to it in order to achieve good results. Usually, this is done by fine-tuning the entire sentence embedding model for the domain of interest. While this approach yields state-of-the-art results, all of the model's weights are updated during fine-tuning, making this method resource-intensive. Therefore, instead of fine-tuning entire sentence embedding models for each target domain individually, we propose to train lightweight adapters. These domain-specific adapters do not require fine-tuning all underlying sentence embedding model parameters. Instead, we only train a small number of additional parameters while keeping the weights of the underlying sentence embedding model fixed. Training domain-specific adapters allows always using the same base model and only exchanging the domain-specific adapters to adapt sentence embeddings to a specific domain. We show that using adapters for parameter-efficient domain adaptation of sentence embeddings yields competitive performance within 1% of a domain-adapted, entirely fine-tuned sentence embedding model while only training approximately 3.6% of the parameters.

1 Introduction

Learning sentence embeddings is an essential task in natural language processing (NLP) and has already been extensively investigated in the literature (Kiros et al., 2015; Hill et al., 2016; Conneau et al., 2017; Logeswaran and Lee, 2018; Cer et al., 2018; Reimers and Gurevych, 2019; Gao et al., 2021; Wu et al., 2022; Schopf et al., 2023d,a). Sentence embeddings are especially useful in information retrieval (Lewis et al., 2020; Schopf et al., 2022;

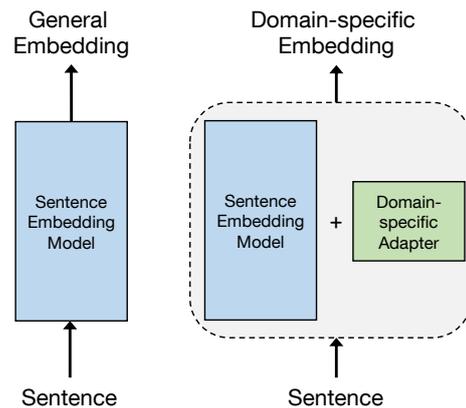


Figure 1: Sentence embedding models are usually trained to obtain state-of-the-art sentence representations for general semantic textual similarity tasks. By injecting domain-specific knowledge of adapters into the sentence embedding model, we can efficiently adapt the resulting representations for semantic textual similarity tasks in different domains.

Schneider et al., 2022) or unsupervised text classification settings (Schopf et al., 2021, 2023b,c). Lately, the most popular approach for sentence embedding learning is to fine-tune pretrained language models with a contrastive learning objective (Liu et al., 2021; Zhang et al., 2022; Chuang et al., 2022; Nishikawa et al., 2022; Cao et al., 2022; Jiang et al., 2022). While this approach provides state-of-the-art results, all of the model's weights are updated during fine-tuning, making this method resource-intensive. This is a problem, particularly when domain-specific models are needed. Then, a specialized model must be trained for each domain of interest, resulting in resource-intensive training.

Recently, *adapters* have emerged as a parameter-efficient strategy to fine-tune Language Models (LMs). Adapters do not require fine-tuning of all parameters of the pretrained model and instead introduce a small number of task-specific parameters while keeping the underlying pretrained lan-

guage model fixed (Pfeiffer et al., 2021a). They enable efficient parameter sharing between tasks and domains by training many task-specific, domain-specific, and language-specific adapters for the same model, which can be exchanged and combined post-hoc (Pfeiffer et al., 2020a). Therefore, many different adapter architectures have been proposed for various domains and tasks (Pfeiffer et al., 2020b, 2021b; Vidoni et al., 2020; He et al., 2021; Le et al., 2021; Parović et al., 2022; Lee et al., 2022). However, to the best of our knowledge, no method currently exists for efficient domain adaptation of sentence embeddings using adapters.

In this paper, we aim to bridge this gap by proposing approaches for adapter-based domain adaptation of sentence embeddings, allowing us to train models for many different domains efficiently. Therefore, we investigate how to adapt general pretrained sentence embedding models to different domains using domain-specific adapters. As shown in Figure 1, this allows always using the same base model to adapt sentence embeddings to a specific domain and only needing to exchange the domain-specific adapters. Accordingly, we train lightweight adapters for each domain and avoid expensive training of entire sentence embedding models.

2 Related Work

Adapters have been introduced by Hounsby et al. (2019) as a parameter-efficient alternative for task-specific fine-tuning of language models. Since their introduction, adapters have been used to fine-tune models for single tasks as well as in multi-task settings (Pfeiffer et al., 2021a). Usually, adapters are used to solve tasks such as classification (Lauscher et al., 2020), machine translation (Baziotis et al., 2022), question answering (Pfeiffer et al., 2022), or reasoning (Pfeiffer et al., 2021a). While there exist adapters for semantic textual similarity (STS) tasks on the *AdapterHub* (Pfeiffer et al., 2020a), these are trained on general STS datasets using a task-unspecific pretrained language model as a basis. We, however, focus on adapting pretrained sentence embedding models to specific domains using adapters.

3 Method

We assume we have a base sentence embedding model from the source domain and labeled datasets for each target domain. Instead of fine-tuning the

entire sentence embedding model for each target domain individually, we train lightweight adapters for each domain. This domain-specific fine-tuning with adapters involves adding a small number of new parameters to the sentence embedding model. During training, the parameters of the sentence embedding model are frozen, and only the weights of the adapters are updated. Formally, we adopt the general definition for adapter-based fine-tuning of Pfeiffer et al. (2021a) as follows:

For each of the N domains, the sentence embedding model is initialized with parameters Θ_0 . Additionally, a set of new and randomly initialized adapter parameters Φ_n are introduced. The parameters Θ_0 are fixed and only the parameters Φ_n are trained. Given training data D_n and a loss function L , the objective for each domain $n \in 1, \dots, N$ is of the form:

$$\Phi_n \leftarrow \underset{\Phi}{\operatorname{argmin}} L(D_n; \Theta_0, \Phi) \quad (1)$$

Usually, the adapter parameters Φ_n are significantly less than the parameters Θ_0 of the base model (Pfeiffer et al., 2021a), e.g., only 3.6% of the parameters of the pretrained model in Hounsby et al. (2019).

4 Experiments

In this section, we describe the used adapter architectures, loss functions, and datasets. In all experiments, we use $\text{SimCSE}_{\text{sup-bert-base}}$ (Gao et al., 2021) as the base sentence embedding model. It is trained on natural language inference (NLI) datasets (Bowman et al., 2015; Williams et al., 2018) for STS tasks in the general domain. We fine-tune all models and adapters for five epochs using a learning rate of $1e^{-5}$.

4.1 Adapter Architectures

We investigate how different adapter architectures affect the domain adaptability of sentence embedding models.

Hounsby-Adapter This adapter, introduced by Hounsby et al. (2019), uses a bottleneck architecture. The adapter modules are added after both the multi-head attention and feed-forward block in each transformer layer (Vaswani et al., 2017) of the base model. The adapter layers transform their input into a very low-dimensional representation and upsample it again to the same dimension in the

output. This generates a parameter-efficient lower-dimensional representation while most information is kept.

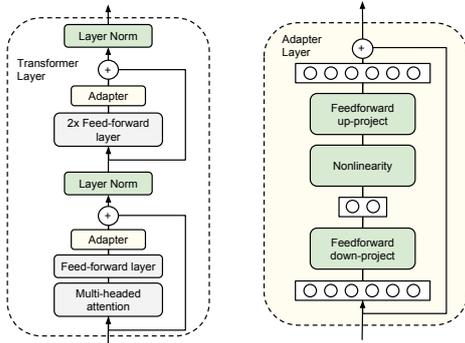


Figure 2: Housby-Adapter architecture as introduced by Housby et al. (2019). On the left side, the adapter is illustrated to be added twice to each transformer layer. Once after the multi-head attention and once after the feed-forward layer. On the right side, the bottleneck architecture of the adapter is presented.

Pfeiffer-Adapter This adapter, introduced by Pfeiffer et al. (2021a), also uses a bottleneck architecture. However, the adapter modules are added only after the feed-forward block in each transformer layer of the base model. This architecture allows merging multiple adapters trained on different tasks. In this work, however, this multitask learning capability is not needed, and we only use the single-task mode.

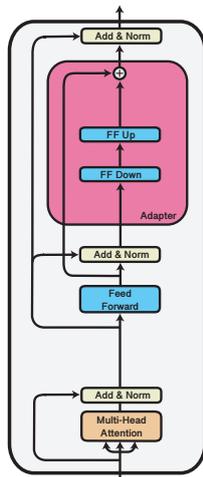


Figure 3: Pfeiffer-Adapter architecture as introduced by Pfeiffer et al. (2021a). Unlike the Housby-Adapter, a single Pfeiffer-Adapter is added in each transformer block only after the forward layer.

K-Adapter This adapter, introduced by Wang et al. (2021), works as outside plug-in for the base model. Each adapter model consists of K adapter

layers containing N transformer layers and two projection layers across which a skip connection is applied. The adapter layers combine the output of an intermediate transformer layer in the base model with the output of a previous adapter layer. To generate the final output, the last hidden states of the adapter are concatenated with the last hidden states of the base model and transformed into the correct output dimension with a simple dense layer.

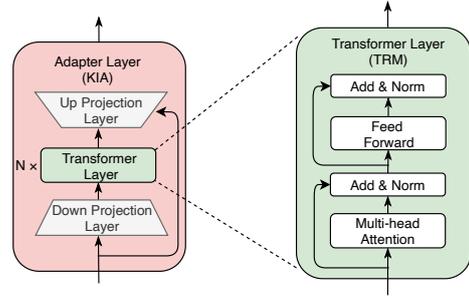


Figure 4: K-Adapter architecture as introduced by Wang et al. (2021). The adapter layer (left) consists of two projection layers, $N = 2$ transformer layers, and a skip connection between two projection layers. The adapter layers are plugged among different transformer layers of the base model. The final output consists of the concatenated last hidden states of the adapter and the base model.

For reference, Table 1 shows the number of parameters per adapter model compared to commonly used base models, highlighting the efficient nature of adapters.

	BERT-base	RoBERTa-large
No. of Parameters Base Model	110M	355M
No. of Parameters Housby-Adapter	4M	6M
No. of Parameters Pfeiffer-Adapter	10M	12M
No. of Parameters K-Adapter	47M	47M

Table 1: Number of trainable Parameters for different base models and adapter architectures.

4.2 Loss Functions

We investigate two different loss functions that are proven to teach models to learn a notion of STS from triplets of examples. We assume a set of triplets $\mathcal{D} = \{(x_i, x_i^+, x_i^-)\}$, where x_i is an anchor sentence, x_i^+ is a positive sample and x_i^- is a negative sample. With $h_i, h_i^+,$ and h_i^- as represen-

Datasets → Models ↓	AskUbuntu	SciDocs				Average
		Cite	CC	CR	CV	
<i>Out-of-the-box</i> SimCSE (<i>lower bound</i>)	60.3	79.3	82.10	76.87	78.36	75.39
ℓ ₁ Houlsby-Adapter	<u>64.0</u>	88.2	88.69	82.42	<u>83.99</u>	81.46
ℓ ₁ Pfeiffer-Adapter	63.8	87.8	<u>88.73</u>	81.65	83.27	81.05
ℓ ₁ K-Adapter	62.5	85.6	87.70	80.09	82.85	79.75
<i>In-domain supervised</i> SimCSE (<i>upper bound</i>)	65.3	88.0	87.74	84.15	83.32	81.70
ℓ ₂ Houlsby-Adapter	64.5	<u>87.3</u>	89.01	<u>82.41</u>	84.42	81.53
ℓ ₂ Pfeiffer-Adapter	64.2	87.0	88.63	81.98	84.41	81.24
ℓ ₂ K-Adapter	62.8	85.3	87.92	80.05	83.29	79.87
<i>In-domain supervised</i> SimCSE (<i>upper bound</i>)	65.2	88.3	88.11	84.46	83.63	81.94

Table 2: Evaluation results of the adapter-based domain adaptation using the different loss functions ℓ_1 and ℓ_2 . The evaluation metric is Mean Average Precision (MAP). We show the performance of the SimCSE model without domain-specific fine-tuning as a lower bound. Additionally, we show the performance of SimCSE models using traditional fine-tuning with the respective loss functions as upper bounds. For the upper bounds, all model weights have been updated during training. In contrast, only the adapter weights were updated during adapter training while the base model parameters were frozen. In bold, we highlight the best adapter performance overall and underline the best adapter results per loss function.

tations of x_i , x_i^+ , and x_i^- , we use the triplet margin loss function of Cohan et al. (2020) as follows:

$$\ell_1 = \max\{(d(h_i, h_i^+) - d(h_i, h_i^-) + m), 0\} \quad (2)$$

where d is the L2 norm distance function and m is the loss margin hyperparameter set to 1.

Additionally, we use the contrastive objective of Gao et al. (2021) as follows:

$$\ell_2 = -\log \frac{e^{\text{sim}(h_i, h_i^+)/\tau}}{\sum_{j=1}^N (e^{\text{sim}(h_i, h_j^+)/\tau} + e^{\text{sim}(h_i, h_j^-)/\tau})} \quad (3)$$

with a mini-batch of N triplets, a temperature hyperparameter τ , which is empirically set to 0.05, and $\text{sim}(h_1, h_2)$ as the cosine similarity $\frac{h_1 \cdot h_2}{\|h_1\| \cdot \|h_2\|}$.

4.3 Data

We use datasets from two different domains to evaluate the domain adaptation abilities of our approach. We randomly split both domain-specific datasets into 90% training and 10% test datasets.

SciDocs The SciDocs dataset (Cohan et al., 2020) consists of scientific papers and their citation information. As model input, we concatenate the titles and abstracts of papers with the [SEP] token. Since our model has a maximum input length of 512 tokens, the input is cut off after this threshold. A

positive sample is defined as a directly referenced paper for each anchor sample. A negative sample is a paper referenced by the positive sample but not by the anchor sample itself. This approach ensures that all samples address the same topic, but the positive sample is more related to the anchor sample than the negative one.

AskUbuntu The AskUbuntu dataset (Lei et al., 2016) consists of user posts from the technical forum AskUbuntu. It already includes sentence pairs that are deemed similar. Therefore, anchor- and positive samples are easily found. Since the dataset inherently consists of sentences about a similar topic, the operating system Ubuntu, negative sentences can easily be retrieved by sampling different sentences. The dataset originates from a technical domain and is quite different from the scientific domain of SciDocs.

5 Evaluation

Table 2 shows the results obtained when adapting sentence embedding models to different domains with adapters. To put the adapter results into perspective, we also evaluate the performance of the SimCSE base model, which is not adapted to the specific domains, as a lower bound. Furthermore, we use traditional domain-specific fine-tuning by training all parameters of the SimCSE base model with the respective loss functions as upper bounds.

The evaluation reveals that adapter-based domain adaptation yields competitive results compared to fine-tuning the entire base model. In particular, the Hounsby and Pfeiffer adapters perform very well with both loss functions, even though they use only a fraction of the parameters of the upper bounds. The slightly larger K-Adapter, however, performs considerably worse than the other adapters investigated. We conclude that the bottleneck architecture is more suitable than the external plug-in architecture for domain adaptation of sentence embedding models. In particular, the Hounsby adapter, although the smallest among the adapters investigated, yields the best results for both loss functions. Using the out-of-the-box SimCSE model without domain adaptation results in considerably worse performance, indicating the overall importance of domain-specific fine-tuning for sentence embedding models.

Furthermore, the contrastive loss function ℓ_2 performs consistently better than ℓ_1 . Our results align with the observations of Gao et al. (2021) who conclude that the contrastive objective ensures a distribution of embeddings around the entire embedding space. In contrast, ℓ_1 may yield learned representations occupying a narrow vector space cone, which severely limits their expressiveness.

From the obtained results, we conclude that using the Hounsby-Adapter architecture together with the contrastive objective ℓ_2 is most suitable for parameter-efficient domain adaptation of sentence embedding models. This adapter approach shows performance that is within 1% of the supervised, entirely fine-tuned SimCSE model, while only training approximately 3.6% of the parameters.

6 Conclusion

In this work, we proposed the use of adapters for parameter-efficient domain adaptation of sentence embedding models. In contrast to fine-tuning the entire sentence embedding model for a particular domain, adapters add a small number of new parameters that are updated during training while the weights of the sentence embedding model are fixed. We showed that adapter-based domain adaptation of sentence embedding models yields competitive results compared to fine-tuning the entire model, although only a fraction of the parameters are trained. In particular, we show that using the Hounsby-Adapter architecture together with a contrastive objective yields promising results for parameter-

efficient domain adaptation of sentence embedding models.

References

- Christos Baziotis, Mikel Artetxe, James Cross, and Shruti Bhosale. 2022. [Multilingual machine translation with hyper-adapters](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1170–1185, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Rui Cao, Yihao Wang, Yuxin Liang, Ling Gao, Jie Zheng, Jie Ren, and Zheng Wang. 2022. [Exploring the impact of negative samples of contrastive learning: A case study of sentence embedding](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3138–3152, Dublin, Ireland. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljagic, Shang-Wen Li, Scott Yih, Yoon Kim, and James Glass. 2022. [DiffCSE: Difference-based contrastive learning for sentence embeddings](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4207–4218, Seattle, United States. Association for Computational Linguistics.
- Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. [SPECTER: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282, Online. Association for Computational Linguistics.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of*

- the 2017 Conference on Empirical Methods in Natural Language Processing, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. **SimCSE: Simple contrastive learning of sentence embeddings**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jiawei Low, Lidong Bing, and Luo Si. 2021. **On the effectiveness of adapter-based tuning for pretrained language model adaptation**. 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 2208–2222, Online. Association for Computational Linguistics.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. **Learning distributed representations of sentences from unlabelled data**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1367–1377, San Diego, California. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. **Parameter-efficient transfer learning for NLP**. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Yuxin Jiang, Linhan Zhang, and Wei Wang. 2022. **Improved universal sentence embeddings with prompt-based contrastive learning and energy-based learning**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3021–3035, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. **Skip-thought vectors**. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Anne Lauscher, Olga Majewska, Leonardo F. R. Ribeiro, Iryna Gurevych, Nikolai Rozanov, and Goran Glavaš. 2020. **Common sense or world knowledge? investigating adapter-based knowledge injection into pretrained transformers**. In *Proceedings of Deep Learning Inside Out (DeeLIO): The First Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 43–49, Online. Association for Computational Linguistics.
- Hang Le, Juan Pino, Changhan Wang, Jiatao Gu, Didier Schwab, and Laurent Besacier. 2021. **Lightweight adapter tuning for multilingual speech translation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 817–824, Online. Association for Computational Linguistics.
- Jaeseong Lee, Seung-won Hwang, and Taesup Kim. 2022. **FAD-X: Fusing adapters for cross-lingual transfer to low-resource languages**. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 57–64, Online only. Association for Computational Linguistics.
- Tao Lei, Hrishikesh Joshi, Regina Barzilay, Tommi Jaakkola, Kateryna Tymoshenko, Alessandro Moschitti, and Lluís Màrquez. 2016. **Semi-supervised question retrieval with gated convolutions**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1279–1289, San Diego, California. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. **Retrieval-augmented generation for knowledge-intensive nlp tasks**. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Che Liu, Rui Wang, Jinghua Liu, Jian Sun, Fei Huang, and Luo Si. 2021. **DialogueCSE: Dialogue-based contrastive learning of sentence embeddings**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2396–2406, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Lajanugen Logeswaran and Honglak Lee. 2018. **An efficient framework for learning sentence representations**. In *International Conference on Learning Representations*.
- Sosuke Nishikawa, Ryokan Ri, Ikuya Yamada, Yoshimasa Tsuruoka, and Isao Echizen. 2022. **EASE: Entity-aware contrastive learning of sentence embedding**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3870–3885, Seattle, United States. Association for Computational Linguistics.
- Marinela Parović, Goran Glavaš, Ivan Vulić, and Anna Korhonen. 2022. **BAD-X: Bilingual adapters improve zero-shot cross-lingual transfer**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational*

- Linguistics: Human Language Technologies*, pages 1791–1799, Seattle, United States. Association for Computational Linguistics.
- Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin Steitz, Stefan Roth, Ivan Vulić, and Iryna Gurevych. 2022. [xGQA: Cross-lingual visual question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2497–2511, Dublin, Ireland. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021a. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2021b. [UNKs everywhere: Adapting multilingual language models to new scripts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10186–10203, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Phillip Schneider, Tim Schopf, Juraj Vladika, Mikhail Galkin, Elena Simperl, and Florian Matthes. 2022. [A decade of knowledge graphs in natural language processing: A survey](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 601–614, Online only. Association for Computational Linguistics.
- Tim Schopf, Karim Arabi, and Florian Matthes. 2023a. [Exploring the landscape of natural language processing research](#).
- Tim Schopf, Daniel Braun, and Florian Matthes. 2021. [Lbl2vec: An embedding-based approach for unsupervised document retrieval on predefined topics](#). In *Proceedings of the 17th International Conference on Web Information Systems and Technologies - WEBIST*, pages 124–132. INSTICC, SciTePress.
- Tim Schopf, Daniel Braun, and Florian Matthes. 2023b. [Evaluating unsupervised text classification: Zero-shot and similarity-based approaches](#). In *Proceedings of the 2022 6th International Conference on Natural Language Processing and Information Retrieval, NLPPIR '22*, page 6–15, New York, NY, USA. Association for Computing Machinery.
- Tim Schopf, Daniel Braun, and Florian Matthes. 2023c. [Semantic label representations with lbl2vec: A similarity-based approach for unsupervised text classification](#). In *Web Information Systems and Technologies*, pages 59–73, Cham. Springer International Publishing.
- Tim Schopf, Emanuel Gerber, Malte Ostendorff, and Florian Matthes. 2023d. [Aspectcse: Sentence embeddings for aspect-based semantic textual similarity using contrastive learning and structured knowledge](#).
- Tim Schopf, Simon Klimek, and Florian Matthes. 2022. [Patternrank: Leveraging pretrained language models and part of speech for unsupervised keyphrase extraction](#). In *Proceedings of the 14th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2022) - KDIR*, pages 243–248. INSTICC, SciTePress.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Marko Vidoni, Ivan Vulić, and Goran Glavas. 2020. [Orthogonal language and task adapters in zero-shot cross-lingual transfer](#). *CoRR*, abs/2012.06460.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2021. [K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1405–1418, Online. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

Xing Wu, Chaochen Gao, Liangjun Zang, Jizhong Han, Zhongyuan Wang, and Songlin Hu. 2022. [ESim-CSE: Enhanced sample building method for contrastive learning of unsupervised sentence embedding](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3898–3907, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Miaoran Zhang, Marius Mosbach, David Adelani, Michael Hedderich, and Dietrich Klakow. 2022. [MCSE: Multimodal contrastive learning of sentence embeddings](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5959–5969, Seattle, United States. Association for Computational Linguistics.