

# CUNI Submission to MRL 2023 Shared Task on Multi-lingual Multi-task Information Retrieval

Jindřich Helcl\* and Jindřich Libovický\*

Charles University, Faculty of Mathematics and Physics  
Institute of Formal and Applied Linguistics  
V Holešovičkách 747/2, 180 00 Prague, Czech Republic  
{helcl, libovicky}@ufal.mff.cuni.cz

## Abstract

We present the Charles University system for the MRL 2023 Shared Task on Multi-lingual Multi-task Information Retrieval. The goal of the shared task was to develop systems for named entity recognition and question answering in several under-represented languages. Our solutions to both subtasks rely on the translate-test approach. We first translate the unlabeled examples into English using a multilingual machine translation model. Then, we run inference on the translated data using a strong task-specific model. Finally, we project the labeled data back into the original language. To keep the inferred tags on the correct positions in the original language, we propose a method based on scoring the candidate positions using a label-sensitive translation model. In both settings, we experiment with finetuning the classification models on the translated data. However, due to a domain mismatch between the development data and the shared task validation and test sets, the finetuned models could not outperform our baselines.

## 1 Introduction

Pre-trained language models reach state-of-the-art results in most current natural language processing (NLP) tasks. Whereas in high-resource languages such as English, we observe in-context learning capabilities and emergent abilities (Wei et al., 2022), in less-resourced languages, the results are more modest (Lai et al., 2023a), mostly due to the lack of necessary data needed to train really large models. Moreover, there is usually not enough task-specific data available in these languages. This leads to attempts to reuse the (high-resource) language model capabilities in other (low-resource) languages.

Most of the proposed methods are either based on transfer learning (Lauscher et al., 2020; Yu and Joty, 2021; Zheng et al., 2021; Schmidt et al., 2022) or machine translation (MT), both during training

and at test time (e.g. mentioned as a baseline by Conneau et al., 2020, 2018).

The MRL 2023 Shared Task on Multi-lingual Multi-task Information Retrieval aims to explore these methods further, applied to many low-resource languages. The participants were tasked to build models for two subtasks: named entity recognition (NER) and question answering (QA).

The shared task setup is inspired by the XTREME-UP dataset (Ruder et al., 2023), which focuses on the most needed tasks for under-resourced languages: gathering data in a digital form (speech recognition, optical character recognition, transliteration) and making information in these languages accessible (NER, QA, retrieval for QA). This dataset contains a relatively small amount of data for multiple tasks on low-resource languages, featuring 88 languages in total, including QA datasets for 4 languages and NER datasets for 20 languages.

The shared task evaluation campaign focused on Igbo, Indonesian (QA only), Alsatian,<sup>1</sup> Turkish, Uzbek (QA only), and Yoruba. Out of these languages, only Indonesian is among the XTREME-UP QA datasets, and only Igbo and Yoruba have available NER task data in the benchmark. Upon releasing the validation data close to the end of the campaign, Azerbaijani was added as a surprise language for evaluation (with no data for QA or NER in XTREME-UP).

This setting left the participants with a choice to either collect external training data for languages not present in the benchmark (which was implicitly discouraged by the inclusion of the surprise language) or to develop language-agnostic systems.

Even though a lot of research effort is invested in developing systems that are inherently multi-lingual, typically based on pre-trained massively multilingual models (Artetxe and Schwenk, 2019; Lauscher et al., 2020; Pfeiffer et al., 2020; Xue

\* The author order was determined by a coin toss.

<sup>1</sup>Mistakenly labeled as Swiss German on the task website.

et al., 2021, inter alia), our submission is based on the translate-test approach that was recently shown to perform better than the community previously thought (Artetxe et al., 2023). We rely on the translation quality of a multi-lingual machine translation (MT) system, combined with the strong performance of pre-trained LLMs in English. The main ideas that are common to our approaches to both subtasks are described in Section 2. The particularities of our models which are specific to the NER and QA subtasks, are presented in Sections 3 and 4, respectively, including our results on those tasks.

Overall, we find that the translate-test approach can be useful in a multilingual setting. Our results do not outperform supervised, language-specific models, but are considerably better than zero-shot approaches.

To maximize reproducibility, we built our systems using an automated end-to-end development pipeline implemented in Snakemake (Köster and Rahmann, 2012); we release the code online.<sup>2</sup>

## 2 Main Ideas

In both tasks, we employ the translate-test approach, which can be summarized in the following three steps: First, we translate unlabeled examples from the task language into English using a multi-lingual MT model. Second, we use a pre-trained LLM to perform the task which assigns the labels to the example. Third, we use a label-aware translation model to project the inferred labels back to the target language.

**Translation into English.** In the first step, we translate the unlabeled data into English. In both subtasks, we use the NLLB-3.3B<sup>3</sup> multilingual MT model (Costa-jussà et al., 2022). We discuss the task-specific data processing details further in Sections 3 and 4.

**Task-specific models.** In each subtask, we apply a RoBERTa-large model, which has been finetuned on the task (Liu et al., 2019). This predicts labels for the English data. For NER, these are BIO-encoded labels, marking the span and type of each named entity in the example. Specifically, the output is a sequence of labels of the same length as the input sentence. For QA, the labels mark a span

in the context representing the answer. This is encoded using two numbers, which denote character offsets in the detokenized version of the context paragraph.

**Translation into the target language.** The translate-test approach is less challenging when the labels are language-independent, which is also the case of both subtasks. However, span labeling tasks (such as NER and QA) require careful handling of the projection of the spans, i.e., we need to find the corresponding spans in the original language.

Our systems adopt the label projection method for cross-lingual transfer, originally meant for the translate-train approach (Chen et al., 2023). The authors of the paper finetune the NLLB model<sup>4</sup> to translate texts containing inserted tags so that the tags generated in the translation mark equivalent parts of the source sentence. In contrast to the original use-case of generating the whole target sentence with tags, we already know the target sentence in the shared task scenario. Therefore, we are only interested in the placement of the tags.

To find the best possible placement of the tags, we propose to use the aforementioned finetuned model as a scorer. We place the tags at all possible positions (subject to minimum/maximum span length constraints) and select the highest-scoring candidate. We then either reconstruct the label sequence (in the case of NER) or extract the appropriate passage from the context (for QA).

## 3 Named Entity Recognition

The goal of the NER subtask was to classify words and phrases into one of four categories: person (PER), organization (ORG), location (LOC), and date (DAT). Since most state-of-the-art NER classifiers (including the one we used) use a richer set of labels, we apply rule-based mapping to reduce the label set to the four categories: geopolitical entities (GPE) and facilities (FAC) are replaced with LOC, time with DAT.

The XTREME-UP benchmark contains two NER datasets, MasakhaNER (Adelani et al., 2021) and MasakhaNER 2 (Adelani et al., 2022), both using texts from local news stories and covering 10 and 20 African languages respectively.

The scheme of the translate-test pipeline for this task is shown in Figure 1.

<sup>2</sup><https://github.com/ufal/mr12023-multilingual-ir-shared-task>

<sup>3</sup><https://huggingface.co/facebook/nllb-200-3.3B-easyproject>

<sup>4</sup><https://huggingface.co/ychenNLP/nllb-200-3.3B-easyproject>




Figure 1: A scheme of the NER translate-test pipeline.

**English NER models.** For experiments with English NER, we use the tNER toolkit (Ushio and Camacho-Collados, 2021), which provides several models for this task. We selected two RoBERTa-based models for experiments, finetuned either on Ontonotes5 (Hovy et al., 2006) or on the CoNLL2003 (Tjong Kim Sang and De Meulder, 2003) dataset. We found that using the model built with Ontonotes5<sup>5</sup> leads to better results than the one finetuned on the CoNLL data.<sup>6</sup> Ontonotes5 contains news stories, television and radio transcripts, and web pages. The CoNLL 2003 data is international news from 1996–1997. This means there is a domain mismatch between the most available named entity recognizers and the MasakhaNER datasets.

**Finetuning.** To overcome the domain mismatch, we finetuned the tNER models using the MasakhaNER data. We translated the MasakhaNER training data into English and performed the span projection the same way as at inference time. The finetuning step serves not only as domain adaptation to news stories from the non-English speaking world but also as an adaptation to texts which have been automatically translated from low-resource languages.

**Results.** Table 1 presents the results on the MasakhaNER 1 dataset. Our translate-test approach significantly outperforms zero-shot transfer from English using the XLM-R (Conneau et al., 2020) and XLM-V (Liang et al., 2023) models; however, there is still a large performance gap between the translate-test approach and supervised in-language training.

The MasakhaNER 2 results are shown in Table 2. Similarly to MasakhaNER 1, our results are strictly worse than supervised training. The second line of the table shows the results of a model trained on MasakhaNER 1 but tested on MasakhaNER 2, which contains ten more languages than the first

<sup>5</sup><https://huggingface.co/tner/roberta-large-ontonotes5>

<sup>6</sup><https://huggingface.co/tner/roberta-large-conll2003>

dataset. The results on these additional languages (shown in boxes) mark zero-shot transfer between African languages. Our translate-test approach via English is better than zero-shot using African languages for 5 of the 10 languages.

When compared to related work, our results (average score 61.3%) without finetuning outperform transfer from English using mDeBERTav3 (He et al., 2023) (average score 55.5%). However, they are worse when compared to the translate-train results reported by Chen et al. (2023) (average score 63.4%) that used additional parallel data with projected labels for training.

On both MasakhaNER benchmarks, the Ontonotes5 model is slightly better than CoNLL 2003. Finetuning (which involves training data of the respective datasets) leads to consistent improvements. On MasakhaNER 2, the finetuned model outperforms Chen et al. (2023); however, the training data setups are not easily comparable.

The results on the shared task validation data are in Table 3. Because of the domain mismatch (the shared task validation data are not local news but rather Wikipedia articles), the original Ontonotes5 model performs better. Based on this observation, we decided to use the pipeline using the original Ontonotes5 model *without* finetuning. We omit Yoruba from calculating the average score because most entities are left without annotation in the data.

## 4 Question Answering

The goal of this task is to find an answer to a given question within a given context. In the generative version of this task, the answer may not be taken from the context directly. Figure 2 shows the question-answering processing pipeline we use in our experiments.

**Data Preprocessing.** The XTREME-UP datasets for QA consist of three fields: The question, the context, and the answer. Since the context might be several sentences long, we apply sentence splitting using wtpsplit (Minixhofer et al., 2023). Since the toolkit does not support Alsatian or Swahili,

	amh	hau	ibo	kin	lug	luo	pcm	swa	wol	yor	AVG
Best baseline (supervised)	78.0	91.5	87.7	77.8	84.7	75.3	90.0	89.5	86.3	83.7	84.4
XLNet-R (zero-shot)	25.1	43.5	11.6	9.4	9.5	8.4	36.8	48.9	5.3	10.0	20.9
XLNet-V (zero-shot)	20.6	35.9	45.9	25.0	48.7	10.4	38.2	44.0	16.7	35.8	32.1
Spacy	59.2	58.3	57.7	48.5	52.6	45.8	9.0	60.0	48.1	47.2	48.6
tNER: ConLL2003	57.3	66.7	72.8	57.0	69.4	49.7	65.8	69.4	53.7	59.3	62.1
tNER: Ontonotes5	60.8	62.8	73.3	60.2	69.9	52.5	74.2	70.1	51.4	57.5	63.3
+ finetuning	61.8	70.0	76.4	65.4	70.2	57.5	77.9	74.5	58.2	59.6	67.2

Table 1: F1 scores on the MasakhaNER 1 dataset.

	bam	bjj	ewe	fon	hau	ibo	kin	lug	mos	nya
Best supervised in paper	82.2	75.2	90.3	82.7	87.4	89.6	87.5	89.6	76.4	92.4
Trained on MasakhaNER 1	50.9	49.8	76.2	57.1	88.7	90.1	87.6	90.0	75.0	80.4
Spacy	38.1	16.8	57.0	39.9	48.1	52.0	55.3	65.7	31.5	53.0
tNER: ConLL2003	49.0	21.9	67.5	51.7	66.3	64.9	60.6	74.8	42.5	66.0
tNER: Ontonotes5	46.8	20.5	67.3	48.8	63.6	63.6	64.6	75.2	39.8	69.0
+ finetuning	<b>60.9</b>	25.9	73.7	53.0	67.0	75.3	65.7	75.2	44.7	72.1

	pcm	sna	swa	tsn	twi	wol	xho	yor	zul	AVG
Best supervised in paper	90.1	96.2	92.7	89.4	81.8	86.8	89.9	89.3	90.6	87.1
Trained on MasakhaNER 1	90.2	42.5	93.1	79.4	57.3	87.0	47.4	89.7	64.3	74.0
Spacy	52.5	60.6	67.4	63.4	53.7	46.5	47.7	42.3	56.2	49.9
tNER: ConLL2003	67.7	69.7	70.8	74.8	67.6	61.9	67.0	52.9	66.0	61.2
tNER: Ontonotes5	72.8	72.4	72.7	73.1	62.0	57.7	67.9	55.6	70.3	61.3
+ finetuning	79.2	<b>81.7</b>	75.1	76.2	<b>68.4</b>	65.6	<b>75.8</b>	60.5	<b>70.4</b>	66.7

Table 2: F1 scores on the MasakhaNER 2 dataset. The numbers in boxes denote zero-shot transfer between African languages (i.e., languages that are in MasakhaNER 2 but not in MasakhaNER 1). Bold numbers are results where our approach is better than the zero-shot transfer between African languages.




Figure 2: A scheme of the QA translate-test pipeline.

	als	aze	tur	yor	AVG
CoNLL 2003	38.4	54.6	48.0	4.4	47.0
Ontonotes5	40.3	62.8	54.7	3.1	52.4
+ finetuned	40.9	62.3	51.3	5.9	51.7

Table 3: Results on the shared task validation data. The average does not include Yoruba.

we use the English variant instead. After sentence splitting, we translate everything into English using the NLLB model. Since NLLB does not support Alsatian, we set the source language to German.

**Answering Questions.** For the extractive question answering task, we use a RoBERTa-based model<sup>7</sup> finetuned for question answering to mark the answer spans in the English context. Once the spans are found, we insert tags into the English sentence. To find the right spans in the original language, we use the tag-preserving NLLB model as a scorer and select the highest-scoring span according to the model.

**No Answer Classification.** Since there are examples with no answer in XTREME-UP, we train a classifier to detect such cases. We again use the QA-tuned RoBERTa, which we finetune on 3 epochs of the translated XTREME-UP data. We set the learning rate to  $10^{-5}$ , weight decay to 0.01, and keep the default values of the rest of the hyper-parameters. The classifier achieves 93% accuracy on the development set. However, because the shared task validation set contains only a very small amount of examples with no answer, we decided not to use this classifier in our submissions.

**In-domain Finetuning.** We also implemented in-domain finetuning of the QA RoBERTa model on the XTREME-UP dataset translated into English. Because the answers are represented as spans within the context, we use the same technique to project the spans onto the English translation of the context as we use in span projecting to the original language. Performing grid search and measuring model performance on the development set, we found a learning rate of  $5 \times 10^{-6}$ , gradient norm of 1, warmup ratio of 0.5, and weight decay of 0.1 are the most suitable hyper-parameters.

**Using Generative Models.** We noticed that the shared task validation data did not actually contain examples of extractive question answering.

<sup>7</sup><https://huggingface.co/deepset/roberta-large-squad2>

Instead, the answers were likely written by a human annotator. Therefore, we decided also to submit a contrastive experiment using a generative model, namely Llama 2 (Touvron et al., 2023).<sup>8</sup> For the generation, we use the prompt "Context: {context} Question: {question} Short answer:". We apply rule-based post-processing to remove potential continuations generated after the answer. Details can be found in the corresponding Snakefile in the code repository.

**Results.** Table 4 shows the results of the shared task validation set. Since there is a considerable domain mismatch between the XTREME-UP dataset and the shared task validation and test sets, we see that the in-domain finetuning does not improve the performance – we, therefore, use the baseline systems as our primary submission. Using the generative model, however, achieves a substantial improvement. Because the task was originally aimed at extractive QA, we decided to submit the generative model as a contrastive experiment.

## 5 Conclusions and Discussion

The research community long overlooked the translate-test approach until recently, when Artetxe et al. (2023) showed that it might outperform both translate-train and cross-lingual transfer with sufficiently strong machine translation systems.

With the increasing number of attempts to use large generative language models in cross-lingual setups, we speculate that the translate-test approach will become an important baseline that might not be easy to cross. Methods that work well with multilingual encoders enforce alignment of the intermediate representation (Wu and Dredze, 2020; Hämmerl et al., 2022; Pfeiffer et al., 2022, inter alia). However, in generative setups, this would lead to undesirable language mixing (Li and Murray, 2023). Generative models are also known not to be consistent across languages (Lai et al., 2023b; Wang et al., 2023). Translate-test does not suffer from either of these disadvantages.

We successfully tested the translate-test method in the shared task setup involving span-labeling tasks. We translated the input into English, performed the task using state-of-the-art English models, and projected the results back to the original language. The main technical challenge is that after labeling the spans in English, we need to find

<sup>8</sup><https://huggingface.co/meta-llama/Llama-2-13b-chat-hf>

	als	aze	ind	tur	uzn	yor	AVG
Roberta-large	8.34	17.87	30.36	14.84	23.13	18.56	18.85
Finetuned	7.88	15.83	30.94	12.64	18.60	19.65	17.59
Llama 2	17.43	31.96	34.61	24.51	30.81	19.80	26.52

Table 4: Question answering results on the shared task validation data (chrF).

the corresponding span in the original text. For that purpose, we used an MT model specifically finetuned to preserve tags encoded as brackets. Furthermore, we finetuned the task-specific models on XTREME-UP data automatically translated into English.

Although the shared task claimed to be based on the XTREME-UP benchmark, the actual shared task data have many different characteristics. Instead of local news outlets, the NER data used Wikipedia text, often on generic topics rather than local ones. The QA validation and test data were abstractive, not extractive. Because of that, our finetuned models performed worse than the original ones. Also, generative QA using LLaMA 2 outperformed our original extractive system.

The final results show that building a translate-test pipeline is a viable approach to both cross-lingual NER and QA.

## Limitations

Both validation and test datasets from the shared task are considerably small, especially for QA, where they contain only around 100 examples per language. This might lead to an unreliable comparison between the submitted systems.

The paper does not contain experimental results that would sufficiently back stronger claims about translate-test approaches. We made decisions that appeared to lead to a good performance in the context of the shared task. However, the paper misses ablations that would reliably show that the span projection method is the best. More importantly, this paper does not compare our results with a strong system based on cross-lingual transfer.

None of the system authors speak the languages in the shared task, and neither is particularly familiar with the culture of the respective language communities. The authors did not check the system outputs for harmful or otherwise inappropriate content.

## Acknowledgments

The work was supported by the Charles University project PRIMUS/23/SCI/023. The work described herein has been using services provided by the LINDAT/CLARIAH-CZ Research Infrastructure (<https://lindat.cz>), supported by the Ministry of Education, Youth and Sports of the Czech Republic (Project No. LM2023062).

## References

- David Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba Alabi, Shamsuddeen Muhammad, Peter Nabende, Cheikh M. Bamba Dione, Andiswa Bukula, Rooweither Mabuya, Bonaventure F. P. Dossou, Blessing Sibanda, Happy Buzaaba, Jonathan Mukiiibi, Godson Kalipe, Derguene Mbaye, Amelia Taylor, Fatoumata Kabore, Chris Chinenye Emezue, Anuoluwapo Aremu, Perez Ogayo, Catherine Gitau, Edwin Munkoh-Buabeng, Victoire Memdjokam Koagne, Allahsera Auguste Tapo, Tebogo Macucwa, Vukosi Mavate, Mboning Tchiaze Elvis, Tajuddeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo Lerato Mokono, Ignatius Ezeani, Chiamaka Chukwunke, Mofetoluwa Oluwaseun Adeyemi, Gilles Quentin Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu, and Dietrich Klakow. 2022. [MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4488–4508, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- David Ifeoluwa Adelani, Jade Abbott, Graham Neubig, Daniel D’souza, Julia Kreutzer, Constantine Lignos, Chester Palen-Michel, Happy Buzaaba, Shruti Rijhwani, Sebastian Ruder, Stephen Mayhew, Israel Abebe Azime, Shamsuddeen H. Muhammad, Chris Chinenye Emezue, Joyce Nakatumba-Nabende, Perez Ogayo, Aremu Anuoluwapo, Catherine Gitau, Derguene Mbaye, Jesujoba Alabi, Seid Muhie Yimam, Tajuddeen Rabiou Gwadabe, Ignatius Ezeani, Rubungo Andre Niyongabo, Jonathan Mukiiibi, Verah Otiende, Iroko Orife, Davis David, Samba Ngom, Tosin Adewumi, Paul Rayson, Mofetoluwa Adeyemi, Gerald Muriuki, Emmanuel Anebi, Chiamaka Chukwunke, Nkiruka Odu, Eric Peter Wairagala, Samuel

- Oyerinde, Clemencia Siro, Tobius Saul Bateesa, Temilola Oloyede, Yvonne Wambui, Victor Akinode, Deborah Nabagereka, Maurice Katusiime, Ayodele Awokoya, Mouhamadane MBOUP, Dibora Gebreyohannes, Henok Tilaye, Kelechi Nwaike, Degaga Wolde, Abdoulaye Faye, Blessing Sibanda, Orevaoghene Ahia, Bonaventure F. P. Dossou, Kelechi Ogueji, Thierno Ibrahima DIOP, Abdoulaye Diallo, Adewale Akinfaderin, Tendai Marengereke, and Salomey Osei. 2021. [MasakhaNER: Named entity recognition for African languages](#). *Transactions of the Association for Computational Linguistics*, 9:1116–1131.
- Mikel Artetxe, Vedanuj Goswami, Shruti Bhosale, Angela Fan, and Luke Zettlemoyer. 2023. Revisiting machine translation for cross-lingual classification. *arXiv preprint arXiv:2305.14240*.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Yang Chen, Chao Jiang, Alan Ritter, and Wei Xu. 2023. [Frustratingly easy label projection for cross-lingual transfer](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5775–5796, Toronto, Canada. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Katharina Hämmel, Jindřich Libovický, and Alexander Fraser. 2022. [Combining static and contextualised multilingual embeddings](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2316–2329, Dublin, Ireland. Association for Computational Linguistics.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTaV3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Johannes Köster and Sven Rahmann. 2012. [Snake-make—a scalable bioinformatics workflow engine](#). *Bioinformatics*, 28(19):2520–2522.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023a. Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning. *arXiv preprint arXiv:2304.05613*.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023b. [Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning](#). *CoRR*, abs/2304.05613.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Tianjian Li and Kenton Murray. 2023. [Why does zero-shot cross-lingual generation fail? an explanation and a solution](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12461–12476, Toronto, Canada. Association for Computational Linguistics.
- Davis Liang, Hila Gonen, Yuning Mao, Rui Hou, Naman Goyal, Marjan Ghazvininejad, Luke Zettlemoyer, and Madian Khabza. 2023. [XLM-V: overcoming the vocabulary bottleneck in multilingual masked language models](#). *CoRR*, abs/2301.10472.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Benjamin Minixhofer, Jonas Pfeiffer, and Ivan Vulić. 2023. [Where’s the point? self-supervised multilingual punctuation-agnostic sentence segmentation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1:*

- Long Papers*), pages 7215–7235, Toronto, Canada. Association for Computational Linguistics.
- Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022. [Lifting the curse of multilinguality by pre-training modular transformers](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3479–3495, Seattle, United States. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. [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.
- Sebastian Ruder, Jonathan H Clark, Alexander Gutkin, Mihir Kale, Min Ma, Massimo Nicosia, Shruti Rijhwani, Parker Riley, Jean-Michel A Sarr, Xinyi Wang, et al. 2023. Xtreme-up: A user-centric scarce-data benchmark for under-represented languages. *arXiv preprint arXiv:2305.11938*.
- Fabian David Schmidt, Ivan Vulić, and Goran Glavaš. 2022. [Don't stop fine-tuning: On training regimes for few-shot cross-lingual transfer with multilingual language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10725–10742, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Asahi Ushio and Jose Camacho-Collados. 2021. [T-NER: An all-round python library for transformer-based named entity recognition](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 53–62, Online. Association for Computational Linguistics.
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, Ai Ti Aw, and Nancy F. Chen. 2023. [SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning](#). *CoRR*, abs/2309.04766.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Shijie Wu and Mark Dredze. 2020. [Do explicit alignments robustly improve multilingual encoders?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4471–4482, Online. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Tao Yu and Shafiq Joty. 2021. [Effective fine-tuning methods for cross-lingual adaptation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8492–8501, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bo Zheng, Li Dong, Shaohan Huang, Wenhui Wang, Zewen Chi, Saksham Singhal, Wanxiang Che, Ting Liu, Xia Song, and Furu Wei. 2021. [Consistency regularization for cross-lingual fine-tuning](#). 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 3403–3417, Online. Association for Computational Linguistics.