

UWB at WASSA-2024 Shared Task 2: Cross-lingual Emotion Detection

Jakub Šmíd*, Pavel Příbáň*, Pavel Král*,†

University of West Bohemia, Faculty of Applied Sciences,

*Department of Computer Science and Engineering,

†NTIS – New Technologies for the Information Society

Univerzitní 2732/8, 301 00 Pilsen, Czech Republic

{jaksmid, pribanp, pkral}@kiv.zcu.cz

<https://nlp.kiv.zcu.cz>

Abstract

This paper presents our system built for the WASSA-2024 Cross-lingual Emotion Detection Shared Task. The task consists of two subtasks: first, to assess an emotion label from six possible classes for a given tweet in one of five languages, and second, to predict words triggering the detected emotions in binary and numerical formats. Our proposed approach revolves around fine-tuning quantized large language models, specifically Orca 2, with low-rank adapters (LoRA) and multilingual Transformer-based models, such as XLM-R and mT5. We enhance performance through machine translation for both subtasks and trigger word switching for the second subtask. The system achieves excellent performance, ranking 1st in numerical trigger words detection, 3rd in binary trigger words detection, and 7th in emotion detection.

1 Introduction

Analyzing emotions in text, including emotion detection and other related tasks, is a well-studied area in natural language processing (NLP). This field has been extensively explored in various SemEval (Mohammad et al., 2018; Chatterjee et al., 2019) and WASSA (Klinger et al., 2018) competitions. The goal of WASSA-2024 Shared Task 2 (Maladry et al., 2024) is to predict specific emotions and identify the words that trigger these emotions. Additionally, this study investigates how emotional information can be transferred across five languages. While the training data is provided in English, the evaluation data includes English, Dutch, Russian, Spanish, and French.

Specifically, the task consists of two subtasks:

- 1) Cross-lingual emotion detection: Predicting emotion from six possible classes (*Love, Joy, Fear, Anger, Sadness, Neutral*) in five target languages.
- 2) Classifying words that express emotions: Identifying words that trigger emotions, with the output

format being either binary (assigning a binary value to each token in the text) or numeric (assigning a numeric value to each token in the text).

Figure 1 shows an example from the dataset for both subtasks. For detailed label descriptions, see the annotation guidelines (Singh, Pranaydeep and Maladry, Aaron and Lefever, Els, 2023).




Figure 1: Example tweet with labels for both subtasks.

The first subtask can be considered a sentence-level classification task, while the second can be treated as a token-level classification task. Both subtasks can also be viewed as text-generation tasks (Raffel et al., 2023). Learning the representation of a given text is essential for solving these problems. Neural networks, including convolutional neural networks (CNNs) (Kim, 2014) and recurrent neural networks (RNNs) (Zaremba et al., 2015), alongside their variations, have been employed for this purpose. However, while these models are effective, they primarily rely on static word embeddings to capture semantic meanings. Consequently, they may struggle with more complex linguistic features such as anaphora and long-term dependencies.

Recent NLP research has shifted towards Transformer-based pre-trained language models (PLMs), such as BERT (Devlin et al., 2019) and T5 (Raffel et al., 2023). These models undergo initial pre-training on extensive datasets to grasp language representation intricacies. Subsequently, they can be fine-tuned on labelled data, capitalizing on the knowledge acquired during pre-training. For cross-lingual tasks, multilingual pre-trained models such as mT5 (Xue et al., 2021) and XLM-RoBERTa (XLM-R) (Conneau et al., 2020) have emerged as standard choices (Hu et al., 2020).

Recently, open-source large language models (LLMs), such as LLaMA 2 (Touvron et al., 2023) and Orca 2 (Mitra et al., 2023), have made significant progress across various NLP tasks. These models show remarkable performance in many zero- and few-shot tasks. Nevertheless, they are primarily pre-trained on English, which often necessitates additional fine-tuning to optimize their performance for other languages (Zhang et al., 2023). However, fine-tuning LLMs on non-specialized consumer GPUs presents challenges due to their large number of parameters. Techniques like QLoRA (Detmers et al., 2023) address this issue by employing a quantized 4-bit frozen backbone LLM with a small set of learnable LoRA weights (Hu et al., 2021).

This paper proposes improving cross-lingual emotion detection by combining a quantized Orca 2 LLM, fine-tuned with LoRA, and machine translation. Additionally, we leverage fine-tuned Transformer-based multilingual language models, such as XLM-R and mT5, for trigger word detection. By incorporating alignment-free translation and trigger word switching, we aim to enhance performance further.¹

2 System Description

We conduct experiments using the dataset provided by Maladry et al. (2024), which contains tweets in five different languages.

2.1 Problem Formulation

For both subtasks, the input is a sentence $x = \{x_i\}_{i=1}^L$ containing L tokens. We denote the parameters of the models as Θ , which includes task-specific parameters \mathbf{W} and \mathbf{b} . Given the sentence-label pairs (x^S, y^S) in the source language S , the aim is to predict labels y^T for the sentence x^T in the target language T .

We formulate the emotion detection subtask as a text-generation problem, which can be modelled, for example, with decoder-only Transformer-based models. The decoder-only model calculates the probability of generating the next token y_t at each step t based on previous outputs y_1, \dots, y_{t-1} as

$$P_{\Theta}(y_t|y_1, \dots, y_{t-1}) = \text{Dec}(y_1, \dots, y_{t-1}), \quad (1)$$

where Dec is the decoder function.

¹The code is publicly available at <https://github.com/biba10/UWBWASSA2024SharedTask2>.

We formulate the second subtask as a token-level classification problem. Given the feature representations $\mathbf{h} = \{\mathbf{h}_i\}_{i=1}^L$ for each token in the sentence obtained by the model, we apply a linear layer on top to get the label distribution for each token x_i as

$$P_{\Theta}(y_i|x_i) = \text{softmax}(\mathbf{W}\mathbf{h}_i + \mathbf{b}), \quad (2)$$

where $y \in \mathcal{Y} = \{0, 1\}$. We select the class with the highest probability. To obtain the numerical values, we extract the logits for class 1, which represents “*is a trigger word*” class, and apply the softmax function to these logits to get numbers between 0 and 1. We consider only the first token of each word for both binary and numerical values.

2.2 Label Projection

Following related work in cross-lingual classification (Hassan et al., 2022; Zhang et al., 2021), we translate the English training set into all non-English target languages using the Google API². This approach significantly expands the training data. The process for the emotion detection task is straightforward: we translate the data and retain the original labels.

Creating translated data for the second subtask is more challenging because the translated text may not perfectly align with the original English text, resulting in different token counts between the two. Previous approaches (Mayhew et al., 2017; Fei et al., 2020) often rely on word alignment tools like fastAlign (Dyer et al., 2013) to map token-label pairs from the source sentence to the targeted one. We adopt an alignment-free label projection proposed by Zhang et al. (2021) to create pseudo-labelled data in the target languages.

First, we mark each trigger word in a sentence with predefined special symbols, such as “[]” or “{ }”, before translating it. We use distinct symbols for a sentence containing multiple trigger words. After translation, we extract the trigger words from the translated sentence using the special symbols. The translation system may occasionally overlook these special symbols, leading to their omission. In such cases, we discard the sentences. This process, illustrated in Figure 2, yields pseudo-labelled sentences in the target language.

Furthermore, we create additional datasets for the second subtask by combining data from different languages, as shown in the lower part of Figure 2. Given an English source sentence x^S and its

²<https://translate.google.com/>




Figure 2: Example of label projection method with trigger words switching (lower part) for English and Spanish language pair.

translation x^T , we switch the trigger words to construct two new bilingual sentences: the first, x^{St} , originates from x^S with trigger words expressed in the target language; the second, x^{Ts} , originates from x^T with trigger words in the source language.

We denote the original English dataset as D_S , the translated dataset into all four non-English languages as D_T , the English source dataset with trigger words in other languages as D_{St} , and the translated dataset with trigger words in English as D_{Ts} .

2.3 LoRA

Fine-tuning LLMs like Orca 2 requires significant computational resources due to the model’s extensive parameter count. To address this challenge, Dettmers et al. (2023) propose to use low-rank adapters (LoRA) (Hu et al., 2021) on top of a quantized backbone model. This method employs a small set of trainable parameters called adapters while keeping the original model frozen, thus reducing memory requirements.

For a pre-trained weight matrix \mathbf{W}_0 , LoRA represents it with a low-rank decomposition as

$$\mathbf{W}_0 + \Delta\mathbf{W} = \mathbf{W}_0 + \mathbf{B}\mathbf{A}, \quad (3)$$

where \mathbf{B} and \mathbf{A} are matrices with much lower dimensions than \mathbf{W}_0 . During fine-tuning, \mathbf{W}_0 is frozen while the weights of \mathbf{A} and \mathbf{B} matrices are updated. Figure 3 shows the concept of LoRA.




Figure 3: The illustrative depiction of parameter-saving LoRA fine-tuning proposed by Hu et al. (2021).

2.4 Models

We fine-tune the large versions of XLM-R (Conneau et al., 2020) and mT5 (Xue et al., 2021), as

well as the 13B version of Orca 2 (Mitra et al., 2023), using the HuggingFace Transformers library³ (Wolf et al., 2020). The Orca 2 model is used for the first subtask, while the XLM-R model and the encoder part of the mT5 model are employed for the second subtask.

2.5 Experimental Setup

For the initial hyperparameter selection, we sample 10% of the English training dataset as validation data for all experiments, following other cross-lingual work (Hu et al., 2019; Zhang et al., 2021) that simulate true unsupervised settings. This 10% is also excluded from the translation process that creates additional datasets. We use the original validation data in five languages to further reduce the number of models.

For both subtasks, we always fine-tune the models on the original English dataset D_S (excluding the 10% of data used for validation) and explore the impact of incorporating additional data. Specifically, for the first subtask, we experiment with adding the translated dataset D_T . For the second subtask, we explore the following options: adding the translated dataset D_T , datasets with switched trigger words (D_{St} and D_{Ts}), or both.

Figure 4 shows the prompt used for fine-tuning the Orca 2 model, where we train the model to output the specific emotion class in a textual format.

Predict one emotion label for the given text. The possible labels are: “Love”, “Joy”, “Anger”, “Fear”, “Sadness”, “Neutral”.
Answer in one following format: “Label: <emotion_label>”

Figure 4: Prompt for the emotion detection task.

2.5.1 Hyperparameters

We employ the AdamW optimizer (Loshchilov and Hutter, 2019) with a batch size 16 for all models.

³<https://github.com/huggingface/transformers>

For Orca 2, we use QLoRA for fine-tuning with 4-bit quantization, setting LoRA $r = 64$ and $\alpha = 16$, a learning rate $2e-4$ with linear decay, and applying LoRA adapters on all linear Transformer block layers. We fine-tune the model for up to 5 epochs. For other models, we fine-tune them for up to 30 epochs and search for the learning rate from $\{2e-6, 2e-5, 5e-5, 2e-4\}$ using a constant scheduler. All experiments are conducted on an NVIDIA A40 with 48 GB GPU memory.

2.5.2 Evaluation Metrics

The main metric for the emotion detection subtask is the macro-averaged F1 score. The primary metric for the binary trigger word detection subtask is the token-level F1 score, calculated on a token level and averaged across instances. A new metric called accumulated precise importance attribution is used for the numerical trigger word detection subtasks. After normalization (ignoring negative values and ensuring the attributions for each sentence add up to 1), this metric sums up the attributions for each trigger word (i.e., the tokens with a label 1).

3 Results

This section presents the results.

3.1 Emotion Detection

Table 1 shows the results of the emotion detection subtask. The results indicate that the Orca 2 largely benefits from the additional data translated into target languages, improving the results by over 3% on test data and by more than 7% on the validation set. The best model achieves a test score of 59.10, ranking seventh in the competition.

Dataset	Dev	Test	Rank
D_S	50.12	55.73	
$D_S + D_T$	57.74	59.10	7.
Baseline		44.76	
Best		62.95	

Table 1: F1 macro scores on the emotion detection task with the Orca 2 model compared to baseline and best results (Maladry et al., 2024). **Bold** indicates the officially announced results and their competition rank.

Figure 5 and Figure 6 show the confusion matrices for the Orca 2 model fine-tuned on English data only and both English and translated data, respectively. The model fine-tuned on translated data demonstrates significantly better performance

for the *Fear* (0.52 vs 0.29) and *Joy* (0.66 vs 0.55) classes while maintaining similar performance for other classes compared to the model fine-tuned only on English data.




Figure 5: Confusion matrix on test data for the Orca 2 model fine-tuned on English data only.

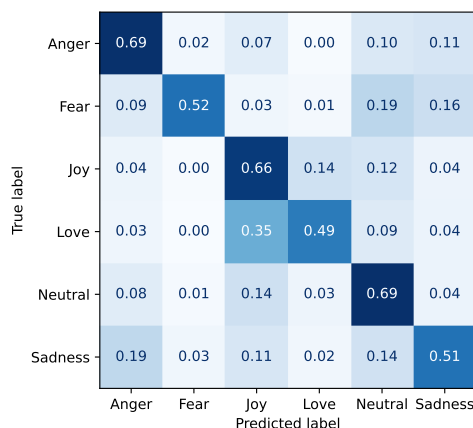


Figure 6: Confusion matrix on test data for the Orca 2 model fine-tuned on English and translated data.

Misclassified labels tend to cluster by sentiment. For instance, the *Love* label is frequently misclassified as *Joy*. The *Neutral*, *Anger*, and *Joy* classes appear to be the easiest to predict, likely due to their higher representation in the training and test data, as shown in Table 2.

Label	Train	Dev	Test
Anger	1,028	129	614
Fear	143	14	77
Joy	1,293	102	433
Love	579	40	190
Neutral	1,397	157	916
Sadness	560	58	270

Table 2: Label distribution for the emotion detection.

3.2 Trigger Words Detection

Table 3 presents the results of the trigger words detection subtask, evaluating the performance for binary and numerical outputs using various training combinations. For binary classification, XLM-R achieves the highest test score (59.19) when trained on the full combination of datasets ($D_S + D_T + D_{St} + D_{Ts}$), ranking third in the competition. Conversely, mT5 excels in numerical value prediction with a top test score of 70.52 under the same training conditions, securing first place. In most cases, the mT5 model’s overall performance on numerical triggers is substantially worse than the performance of XLM-R. However, the results improve significantly when all dataset combinations are used.

Determining the best dataset combination and model is challenging due to the similar results achieved across many cases. Variability from random seeds during fine-tuning can obscure slight performance differences, which may not truly indicate superiority but rather random fluctuations. Nonetheless, our best test set results are obtained by training on a combination of the original, translated, and trigger words switched datasets. Combining all datasets for numerical trigger word detection using the mT5 model significantly improves the second-best settings by 8%.

Model	Dataset	Binary			Numerical		
		Dev	Test	Rank	Dev	Test	Rank
XLM-R	D_S	61.85	58.59		70.07	70.06	
	$D_S + D_T$	60.82	56.69		71.52	66.20	
	$D_S + D_{St} + D_{Ts}$	57.97	53.18		73.14	70.16	
	$D_S + D_T + D_{St} + D_{Ts}$	60.46	59.19	3.	70.05	70.02	
mT5	D_S	59.99	58.12		60.11	60.00	
	$D_S + D_T$	62.12	58.06		63.14	62.06	
	$D_S + D_{St} + D_{Ts}$	53.32	48.61		66.18	62.86	
	$D_S + D_T + D_{St} + D_{Ts}$	59.22	56.79		70.92	70.52	1.
Baseline		23.49			21.60		
Best		61.58			70.52		

Table 3: Token-level F1 score for binary trigger words detection and accumulate precise importance for numerical trigger words detection compared to baseline and best results (Maladry et al., 2024). **Bold** indicates the officially announced results and their competition rank.

3.3 Discussion

Overall, the Orca 2 benefits more from the translated data than the XLM-R and mT5 models, likely because the Orca 2 is pre-trained mostly on English data. In contrast, the other two models are multilingual. The translated data adds more knowledge to a model pre-trained mainly on English data than those already exposed to multiple languages. In

addition, trigger word detection tasks may be more prone to translation errors, which could diminish the benefits of translated data for these tasks. However, the mT5 model shows marginal improvements from the translated and trigger words switched datasets for the numerical trigger words detection, improving the results by 8% and achieving the best result overall in the competition.

4 Conclusion

This paper describes our system for the WASSA-2024 Cross-lingual Emotion Detection Shared Task. We propose fine-tuning a quantized large language model with low-rank adapters combined with machine translation for the emotion detection subtask and fine-tuning multilingual Transformer-based models enhanced with machine translation and trigger word switching for the trigger words detection subtask. We show that additional translated data improves the performance. Our system achieves excellent results and ranks first in numerical trigger word detection, third in binary trigger word detection, and seventh in emotion detection.

Acknowledgements

This work has been partly supported by the OP JAC project DigiTech no. CZ.02.01.01/00/23_021/0008402 and by the Grant No. SGS-2022-016 Advanced methods of data processing and analysis. Computational resources were provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

Limitations

The method relies on machine translation, thus its effectiveness depends on translation quality. In addition, we solve each subtask independently. Future research could address solving both subtasks simultaneously, potentially leading to a more robust model that better explains its decisions.

Ethics Statement

We ensure fair and honest analysis while conducting our work ethically and without harming anybody. However, we acknowledge that the language models used in this paper may introduce unintended biases related to race or gender due to pre-training on large corpora.

References

- Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. [SemEval-2019 task 3: EmoContext contextual emotion detection in text](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 39–48, Minneapolis, Minnesota, USA. 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.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). *Preprint*, arXiv:2305.14314.
- 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.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Hao Fei, Meishan Zhang, and Donghong Ji. 2020. [Cross-lingual semantic role labeling with high-quality translated training corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7014–7026, Online. Association for Computational Linguistics.
- Sabit Hassan, Shaden Shaar, and Kareem Darwish. 2022. [Cross-lingual emotion detection](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6948–6958, Marseille, France. European Language Resources Association.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Minghao Hu, Yuxing Peng, Zhen Huang, Dongsheng Li, and Yiwei Lv. 2019. [Open-domain targeted sentiment analysis via span-based extraction and classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 537–546, Florence, Italy. Association for Computational Linguistics.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Roman Klinger, Orphée De Clercq, Saif Mohammad, and Alexandra Balahur. 2018. [IEST: WASSA-2018 implicit emotions shared task](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 31–42, Brussels, Belgium. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.
- Aaron Maladry, Pranaydeep Singh, and Els Lefever. 2024. [Findings of the wassa 2024 exalt shared task on explainability for cross-lingual emotion in tweets](#). In *Proceedings of the 14th Workshop of on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis@ACL 2024*, Bangkok, Thailand.
- Stephen Mayhew, Chen-Tse Tsai, and Dan Roth. 2017. [Cheap translation for cross-lingual named entity recognition](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2536–2545, Copenhagen, Denmark. Association for Computational Linguistics.
- Arindam Mitra, Luciano Del Corro, Shweti Mahajan, Andres Coda, Clarisse Simoes, Sahaj Agarwal, Xuxi Chen, Anastasia Razdaibiedina, Erik Jones, Kriti Agarwal, Hamid Palangi, Guoqing Zheng, Corby Rosset, Hamed Khanpour, and Ahmed Awadallah. 2023. [Orca 2: Teaching small language models how to reason](#). *Preprint*, arXiv:2311.11045.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. [SemEval-2018 task 1: Affect in tweets](#). In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Preprint*, arXiv:1910.10683.

- Singh, Pranaydeep and Maladry, Aaron and Lefever, Els. 2023. Annotation guidelines for labeling emotion in multilingual tweets.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal 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](#). *Preprint*, arXiv:2307.09288.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- 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.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2015. [Recurrent neural network regularization](#). *Preprint*, arXiv:1409.2329.
- Wenxuan Zhang, Ruidan He, Haiyun Peng, Lidong Bing, and Wai Lam. 2021. [Cross-lingual aspect-based sentiment analysis with aspect term code-switching](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9220–9230, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023. [Don't trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7915–7927, Singapore. Association for Computational Linguistics.