

CSIRO-LT at SemEval-2025 Task 11: Adapting LLMs for Emotion Recognition for Multiple Languages

Jiyu Chen^{1*}, Necva Bölücü^{1*}, Sarvnaz Karimi¹, Diego Mollá^{2,1}, Cécile Paris^{1,2}

¹CSIRO Data61, Australia

²Macquarie University, Australia

firstname.lastname@csiro.au diego.molla-ali@mq.edu.au

* indicates co-first authors.

Abstract

Detecting emotions across different languages is challenging due to the varied and culturally nuanced ways of emotional expressions. The *SemEval 2025 Task 11: Bridging the Gap in Text-Based emotion* shared task was organised to investigate emotion recognition across different languages. The goal of the task is to implement an emotion recogniser that can identify the basic emotional states that general third-party observers would attribute to an author based on their written text snippet, along with the intensity of those emotions. We report our investigation of various task-adaptation strategies for LLMs in emotion recognition. We show that the most effective method for this task is to fine-tune a pre-trained multilingual LLM with LoRA setting separately for each language.

1 Introduction

Text-based emotion recognition plays a crucial role in studies related to mental health (Golder and Macy, 2011), emotional intelligence (Turcan et al., 2021), and human-computer interaction (Li et al., 2022). However, recognising emotions in text remains a significant challenge, especially across different languages (Mohammad et al., 2018). Linguistic variations, cultural differences in emotional expression, and the scarcity of annotated data for low-resourced languages make emotion recognition particularly complex (Barrett et al., 2011; Lindquist and Gendron, 2013; Schröder et al., 2013).

The *SemEval-2025 Task 11: Bridging the Gap in Text-Based Emotion* shared task (Muhammad et al., 2025b)¹ is part of the International Workshop on Semantic Evaluation (SemEval). Its objective is to detect *collectively perceived emotions* within text

snippets written in 32 different languages. Collectively perceived emotion refers to the basic emotional states—such as anger, joy, and disgust—that third-party observers from the general public can attribute to a text snippet generated by a writer. Unlike tasks aimed at identifying the writer’s actual emotional states or the emotional states evoked in individual reader (Mohammad, 2022, 2023), this task emphasises the shared perception of emotions by general readers. This distinction is crucial, as perceived emotions can vary significantly from both intended and personally experienced emotions.

The shared task consists of three tracks: (A) *Multi-label Emotion Detection*, where the goal is to predict the perceived emotional states expressed in a given text snippet, including joy, sadness, fear, anger, surprise, and disgust; (B) *Emotion Intensity*, where the objective is to predict the intensity scale of each perceived emotional state for a given text snippet. Each emotional state is rated on a 4-point categorical scale: 0 (no emotion), 1 (low intensity), 2 (moderate intensity), and 3 (high intensity); and (C) *Cross-lingual Emotion Detection*, where the goal is to predict the perceived emotion on text snippets written in a language different from the one used for model training, such as training on English-written data but making emotion recognition on Japanese-written data.

Our team participated in Tracks A and B, focusing on fine-tuning Large Language Models (LLMs). LLMs have contributed to significant improvements across different NLP tasks (Brown, 2020; Touvron et al., 2023). However, previous studies suggest that they are ineffective for emotion classification in zero-shot and few-shot settings, even when provided with In-Context Learning (ICL) prompts (Liu et al., 2024). A potential improvement can be achieved by fine-tuning LLMs with instructions (Zhang et al., 2023; Liu et al., 2024). Drawing inspiration from Liu et al. (2024), we ex-

¹<https://github.com/emotion-analysis-project/SemEval2025-task11>

plore instruction-tuning and continual fine-tuning of multilingual LLMs. We then compare the effectiveness of this adaptation for emotion recognition across different languages by fine-tuning a multilingual model individually for each language. We also propose several adaptation approaches and conduct a comparative analysis across these approaches, specifically for Track A. We develop an adaptation strategy for LLMs in emotion recognition that is effective across languages of varying resourced levels (see the categorisation of resource abundance by Joshi et al. (2020); Üstün et al. (2024)).

2 Methods

2.1 Track A: Multi-label Emotion Detection

We formulate this task as a **binary classification** problem for the detection of each emotional state. For each input, the LLMs predict the occurrence of the six emotions—anger, disgust, fear, joy, sadness, and surprise—independently, determining whether a specific emotional state is *present* (1) or *absent* (0). The final results are aggregated for the input, to represent a multi-label emotion recognition setting. To address class imbalance, we apply oversampling to balance the binary classification instances.

The proposed adaptation strategies are described below:

- **Few-shot:** We apply ICL based on the BM25 (Robertson et al., 1995) scores to retrieve instances from the training set to construct a prompt for the prediction of each test instance (see the prompt template in Section 3.3). That is, we use BM25 scores to rank semantic (bag-of-words) relevancy between the given test instance (as a query) and each training instance (as a document) (Schutze et al., 2008; Robertson et al., 2009). We retrieve the top k most relevant instances and their manually annotated emotional state label as k -shot examples to prompt-tune LLMs for each test instance.
- **Supervised Fine-Tuning (SFT):** We fine-tune pre-trained and instruction-tuned multilingual LLMs using supervised and parameter-efficient settings.
- **English-bridged Adaptation (E-Bridge):** We apply SFT to LLMs on English-written instances and then apply continual SFT for the adaptation to other languages.

- **Marginalisation:** We first apply SFT to LLM on Track B (see details of SFT in Section 2.2), and then use fine-tuned LLM to make predictions for instances on Track A. To align with the binary classification in Track A, predictions of 1–3 are marginalised into 1, indicating the presence of an emotional state, while predictions of zero are retained to represent its absence.

2.2 Track B: Emotion Intensity Detection

We frame this task as a **multi-class classification** problem for each emotion. Given an input text, the LLM predicts whether a specific emotion can be perceived at a given intensity level. The intensity is measured on a four-point categorical scale: 0 for no emotion, 1 for low intensity, 2 for moderate intensity, and 3 for high intensity. Thus, for a text with six emotional states, the LLM processes the input six times, once for each emotion. The final results are aggregated to the input text for the completion of multi-label emotion intensity recognition set by Track B.

To achieve this, we convert each text with multiple emotion intensities into separate instances, one for each emotion. We fine-tune pre-trained and instruction-tuned LLMs on this transformed dataset, using supervised learning to predict the emotional intensity for each emotion independently (SFT) and apply zero-shot, where we only use instruction as a baseline.

3 Experimental Setup

3.1 Shared Task Dataset

The dataset used in the shared task is a combination of EthioEmo (Belay et al., 2025) and BRIGHTER (Muhammad et al., 2025a), which includes emotion annotations for multiple languages. Specifically, 28 languages for Track A and 11 languages for Track B. The statistical details of the annotated languages are detailed by Belay et al. (2025) for Amharic, Oromo, Somali, and Tigrinya, and by Muhammad et al. (2025a) for the remaining languages.

Instruction-tuning on External Dataset: We leveraged an external dataset to instruction-tune LLMs (Liu et al., 2024). This helps the model generalise on the external task by learning how to follow instructions for emotion recognition before fine-tuning on the dataset specific to this shared

Languages	Track A (F_1)			Track B (r)		
	Baseline	Ours	Model	Baseline	Ours	Model
Afrikaans (afr)	37.14	31.61	aya-101	—	—	—
Algerian Arabic (arq)	41.41	52.55	aya-32	1.64	52.11	aya-32
Amharic (amh)	63.83	58.81	aya-32	50.79	15.40	llama
Chinese (chn)	53.08	57.84	aya-32	40.53	54.92	aya-32
Emakhuwa (vmw)	12.14	17.27	aya-101	—	—	—
English (eng)	70.83	77.62	aya-32	64.15	72.09	llama
German (deu)	64.23	65.72	aya-32	56.21	60.98	aya-32
Hausa (hau)	59.55	54.25	aya-101	27.03	37.15	llama
Hindi (hin)	85.51	73.16	aya-32	—	—	—
Igbo (ibo)	47.90	32.56	aya-101	—	—	—
Indonesian (ind)	—	—	—	—	—	—
isiXhosa (xho)	—	—	—	—	—	—
isiZulu (zul)	—	—	—	—	—	—
Javanese (jav)	—	—	—	—	—	—
Kinyarwanda (kin)	46.29	42.95	aya-101	—	—	—
Marathi (mar)	82.22	75.75	aya-101	—	—	—
Moroccan Arabic (ary)	47.16	—	—	—	—	—
Nigerian-Pidgin (pcm)	55.50	21.81	aya-101	—	—	—
Oromo (orm)	12.63	39.24	aya-101	—	—	—
Portuguese (Brazil) (ptbr)	42.57	55.54	aya-32	29.74	48.88	emo-aya
Portuguese (Mozambican) (ptmz)	45.91	—	—	—	—	—
Romanian (ron)	76.23	72.24	aya-101	55.66	59.22	aya-32
Russian (rus)	83.44	89.10	aya-101	87.66	79.26	llama
Somali (som)	45.93	43.28	aya-101	—	—	—
Spanish (Latin American) (esp)	77.44	82.00	aya-101	72.59	69.66	llama
Sundanese (sun)	37.31	48.75	aya-101	—	—	—
Swahili (swa)	22.65	30.31	aya-101	—	—	—
Swedish (swe)	51.98	48.88	aya-101	—	—	—
Tatar (tat)	53.94	53.84	aya-101	—	—	—
Tigrinya (tir)	46.28	49.95	aya-101	—	—	—
Ukrainian (ukr)	53.45	66.40	aya-32	39.94	49.42	emo-aya*
Yoruba (yor)	9.22	29.96	aya-101	—	—	—

Table 1: Effectiveness of baseline and our approaches (SFT or zero-shot (*) of a specific model) on the test set. The baseline results are provided by the task organisers [Muhammad et al. \(2025a\)](#). Note that aya-32 denotes the aya-32b-expanse model, llama denotes the Llama3.1-8B-Instruct model, emo-aya denotes instruction-tuned aya-32b-expanse.

task. The selected external dataset is an extension to the *SemEval-2018 Task 1: Affect in Tweets*, which includes a series of subtasks related to affectual state inference: (1) emotion intensity regression; (2) emotion intensity ordinal classification; (3) valence (sentiment) regression; (4) valence ordinal classification; and, (5) emotion classification ([Mohammad and Kiritchenko, 2018](#); [Mohammad et al., 2018](#)) which overlaps with the tracks of this shared task.

3.2 Evaluation Metrics

For both Track A and Track B, we follow the evaluation metrics specified by the shared task organ-

isers. Track A uses the unweighted average of all per-emotion F_1 scores. Track B uses the average of per-emotion Pearson’s correlation coefficient (r).

3.3 Hyperparameters

LLMs: We use three open-source multilingual LLMs: AYA (aya-101 ([Üstün et al., 2024](#)), aya-32b-expanse ([Dang et al., 2024](#))), and LLAMA (Llama3.1-8B-Instruct ([Dubey et al., 2024](#))). aya-101 (mT5-based) offers broad multilingual support with a smaller size, while aya-32b-expanse (GPT-style) provides larger capacity and similarly wide language coverage. Llama3.1-8B-Instruct (GPT-style) is smaller

Language	Supervised Fine-tuning (SFT)			Marginalisation				Few-shots
	AYA _{ft}	EMO-AYA _{ft}	E-Bridge	AYA _{pt_b}	EMO-AYA _{pt_b}	AYA _{ft_b}	EMO-AYA _{ft_b}	1-shot
English	80.12	73.31	—	64.55	57.94	68.47	72.88	52.31
German	62.31	52.43	56.21	53.92	46.65	60.69	54.56	30.24
Portuguese	61.29	53.42	57.33	44.55	40.56	41.87	45.93	34.11
Russian	89.11	87.13	86.16	54.12	60.56	72.42	55.34	75.01

Table 2: Macro F_1 scores on the development split of Track A for four languages. AYA_{ft} denotes direct SFT of the aya-32b-expanse model on the training set. EMO-AYA_{ft} indicates first performing instruction-tuning on the aya model on the external dataset, followed by SFT on the training set. E-Bridge refers to SFT of the aya model on the training set in English, followed by continual SFT on the remaining three languages. AYA_{pt_b} involves prompt-tuning a model in a zero-shot setting to complete the Track B objective, followed by marginalisation. EMO-AYA_{pt_b} represents applying instruction-tuning to aya model on the external dataset and then prompt-tuning it in a zero-shot setting on Track B, followed by marginalisation. The best results are **boldfaced**.

but supports only seven languages. Due to its broader coverage and capacity, aya-32b-expanse was preferred, with aya-101 as a lightweight alternative. For languages jointly supported by both aya-32b-expanse and Llama3.1-8B-Instruct, such as English and German, model selection also accounts for differences in parameter size.

We employ Low-Rank Adaptation (LoRA) (Hu et al., 2022) and apply 4-bits quantisation (Jacob et al., 2018) to LLMs for parameter-efficient SFT. We set the LoRA rank and alpha parameters to 32 and 64, respectively. The dropout ratio is set to 0.05. We limit both the input source length and the target length to 512. The training epoch size is 10 and the batch size is 2. The learning rate is set to $2e - 5$ for Track A and $5e - 5$ for Track B.

The formulations of instructions for zero-shot, few-shot (ICL), and SFT settings are as below:

- **Track A:** “You are detecting emotions on a statement written in {language}. Statement: {text}. Does this statement express {emotion}? Answer 1 for yes and 0 for no.”
- **Track B:** “Task: Categorize the tweet into an intensity level of the specified emotion E, representing the mental state of the tweeter. 0: no E can be inferred. 1: low amount of E can be inferred. 2: moderate amount of E can be inferred. 3: high amount of E can be inferred. Tweet: {text} Emotion {emotion} Intensity class:”

where {text} is the text content of each instance, and {emotion} is one of the six emotional states (or five, excluding disgust for English and Surprise for Afrikaans). The instruction of Track B is adapted from Liu et al. (2024).

For BM25-based few-shot prompting, we use the rank_bm25 library (Stuart, 2022) and choose the default parameter setting, $b = 0.75$ and $k_1 = 1.5$, for BM25.

We use NVIDIA H100 GPUs running on one node for this experiment.

4 Experimental Results

Results on Testing Set

We present the results of the submitted predictions for ranking (testing) in Table 1, including the baseline results per language provided by Muhammad et al. (2025a). We submitted results for 26 out of 32 languages in Track A and the 11 (all) languages provided in Track B.

We observed noticeable variations in effectiveness across languages for both Track A (macro F_1) and Track B (average r). The baseline approach is mostly effective in higher-resourced languages, such as German, English, and Russian. However, applying instruction-tuning or SFT to LLMs on mid-resourced and lower-resourced language is more effective than the baseline for emotion recognition.

Additionally, model selection plays a crucial role, as larger GPT-based models like aya-32b-expanse outperform smaller mT5-based models like aya-101, particularly in lower-resourced languages where the latter struggles.

We applied the adaptation strategies for the test set prediction based on the highest F_1 and correlation coefficient r achieved on the development set in Track A & B, respectively.

Results on Track A Development Set

We compared the effectiveness of the various adaptation strategies with experiments in English, Ger-

Language	Zero-shot				SFT			
	AYA	EMO-AYA	LLAMA	EMO-LLAMA	AYA _{ft}	EMO-AYA _{ft}	LLAMA _{ft}	EMO-LLAMA _{ft}
Algerian Arabic (arq)	41.55	46.39	16.65	37.44	64.66	9.11	30.61	28.56
Amharic (amh)	7.63	7.37	11.93	16.40	—	—	25.30	24.12
Chinese (chn)	47.61	49.00	35.57	43.01	62.37	48.37	39.83	44.15
English (eng)	57.80	56.16	43.87	57.32	73.81	74.04	75.87	75.23
German (deu)	53.00	43.39	41.61	45.11	55.93	50.04	41.09	44.16
Hausa (hau)	15.08	16.77	16.15	15.34	21.31	16.18	45.13	42.56
Portuguese (Brazilian) (ptbr)	46.71	49.02	31.17	30.96	45.63	53.42	—	—
Romanian (ron)	57.92	47.36	47.98	47.74	63.35	61.22	50.76	52.34
Russian (rus)	56.01	62.34	34.98	50.47	75.44	70.63	83.62	82.10
Spanish (Latin American) (esp)	57.77	54.77	45.80	48.69	68.50	64.36	70.07	69.10
Ukrainian (ukr)	45.68	50.07	24.12	33.05	—	—	41.99	40.56

Table 3: The average r on the development set of Track B. AYA and LLAMA refer to the base models, aya-expense-32b and Llama3.1-8B-Instruct, respectively. EMO-AYA and EMO-LLAMA are their instruction-tuned versions. *_{ft} are fine-tuned versions. The best results are **boldfaced**.

man, Portuguese, and Russian. All comparisons are statistically validated using hypothesis testing with a significance threshold of $p < 0.05$. We observed that directly applying SFT to LLMs on the training set (AYA_{ft}) consistently achieves the highest macro F_1 scores across all four languages, outperforming all of the other experimented adaptation approaches, such as: (i) BM25-based ICL (1-shot), (ii) instruction-tuning on the external dataset before SFT on the training set (EMO-AYA), (iii) bridging the adaptation with English (E-Bridge), or (iv) applying marginalisation to predictions of LLMs fine-tuned on Track B (Table 2).

These results suggest that LLMs may not significantly benefit from instruction tuning, as direct SFT shows greater effectiveness across all four languages. E-Bridge can be viewed as a specialised form of instruction tuning, where the LLM first learns instructions in English instances before being fine-tuned in other languages. This method proves effective for German and Portuguese but is less effective for Russian, possibly due to the closer cultural alignment of German and Portuguese speakers with English speakers (Rinke and Flores, 2021; Wikipedia, 2025).

Results on Track B Development Set

The emotion intensity results across multiple languages using both base and instruction-tuned versions of AYA and LLAMA in zero-shot and SFT settings are shown in Table 3. Fine-tuning significantly improves performance across most languages, demonstrating that task-specific adaptation benefits intensity detection. While fine-tuning offers substantial gains across the languages, the benefits are often more pronounced in higher-resourced languages (i.e., English, Spanish). While AYA

generally demonstrates stronger zero-shot performance, LLAMA benefits more from instruction tuning, showing significant improvements after fine-tuning. Instruction-tuned models (EMO-AYA and EMO-LLAMA) provide some advantages in zero-shot settings, particularly in languages with less training data, but their impact diminishes after fine-tuning, suggesting that instruction-tuning alone is often sufficient for intensity detection.

Higher-resourced languages, such as English, Russian, and Spanish, consistently achieve better results, with both zero-shot and fine-tuned models performing reliably. Mid-resourced languages, including German, Portuguese, and Romanian, show moderate performance, benefiting from fine-tuning but still exhibiting variability depending on the model. The performance of all languages improves with fine-tuning, but challenges persist in making significant gains for languages where the models initially perform poorly. Despite this, the advancements show that fine-tuning and instruction tuning can help optimise model behaviour across languages, and targeted adaptation strategies may further boost results for emotion intensity detection tasks.

5 Conclusions

We participated in the *SemEval 2025 Task 11: Bridging the Gap in Text-Based Emotion* shared task, which included tracks for multi-label emotion detection (Track A) and emotion intensity detection (Track B). For Track A, we approached it as a binary classification problem for each emotion—determining whether an emotional state is present in or absent from a given input text snippet. We found that direct supervised fine-tuning can effectively adapt LLMs for the detection of emotions

for most languages, except for the lower-resourced languages where few-shot learning is more effective. For Track B, we achieved the best results by employing different LLMs (both direct SFT and instruction-tuned) for each language. The results in both tracks suggest that, when adapting LLMs for emotion recognition on most mid-resourced and higher-resourced language, instruction-tuning was not as effective as in other NLP tasks. A more suitable approach is to directly apply supervised fine-tuning of LLMs on task-specific datasets.

Limitations

The exploration of various adaptation strategies was limited to four languages (English, German, Russian, and Portuguese), which may not generalise to other languages, particularly lower-resourced ones or those with different linguistic structures. The models used may reflect biases from the training data, which could affect performance in low-resourced languages. We only explored prompt-tuning and instruction-tuning with the parameter-efficient LoRA setting for adapting LLMs.

Ethical Considerations

We relied on the dataset providers to remove any material from the dataset that may reveal anyone’s identity in their posts used in this study. We guarantee that datasets are only used for scientific or research purposes and are not redistributed or shared with third parties. This project is subject to the ethics approval and agreement provided by the SemEval-2025 task organisers.

References

Lisa Feldman Barrett, Batja Mesquita, and Maria Gendron. 2011. Context in emotion perception. *Current Directions in Psychological Science*, 20(5):286–290.

Tadesse Destaw Belay, Israel Abebe Azime, Abinew Ali Ayele, Grigori Sidorov, Dietrich Klakow, Philip Slusallek, Olga Kolesnikova, and Seid Muhie Yimam. 2025. Evaluating the Capabilities of Large Language Models for Multi-label Emotion Understanding. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3523–3540. Association for Computational Linguistics.

Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith,

Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, Acyr Locatelli, Hangyu Lin, Dwarak Talupuru, Bharat Venkitesh, David Cairuz, Bowen Yang, Tim Chung, Wei-Yin Ko, Sylvie Shang Shi, Amir Shukayev, Sammie Bae, Aleksandra Piktus, Roman Castagné, Felipe Cruz-Salinas, Eddie Kim, Lucas Crawhall-Stein, Adrien Morisot, Sudip Roy, Phil Blunsom, Ivan Zhang, Aidan Gomez, Nick Frosst, Marzieh Fadaee, Beyza Ermis, Ahmet Üstün, and Sara Hooker. 2024. *Aya expand: Combining Research Breakthroughs for a New Multilingual Frontier*. Preprint, arXiv:2412.04261.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Scott A Golder and Michael W Macy. 2011. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. *LoRA: Low-rank adaptation of large language models*. In *International Conference on Learning Representations*.

Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713.

Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293.

Qintong Li, Piji Li, Zhaochun Ren, Pengjie Ren, and Zhumin Chen. 2022. Knowledge bridging for empathetic dialogue generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10, pages 10993–11001.

Kristen A Lindquist and Maria Gendron. 2013. What’s in a word? Language constructs emotion perception. *Emotion Review*, 5(1):66–71.

Zhiwei Liu, Kailai Yang, Qianqian Xie, Tianlin Zhang, and Sophia Ananiadou. 2024. Emollms: A series of emotional large language models and annotation tools for comprehensive affective analysis. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5487–5496.

- Saif Mohammad. 2023. Best Practices in the Creation and Use of Emotion Lexicons. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1825–1836.
- 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.
- Saif Mohammad and Svetlana Kiritchenko. 2018. [Understanding Emotions: A Dataset of Tweets to Study Interactions between Affect Categories](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 198–209. European Language Resources Association (ELRA).
- Saif M Mohammad. 2022. Ethics sheet for automatic emotion recognition and sentiment analysis. *Computational Linguistics*, 48(2):239–278.
- Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Jan Philip Wahle, Terry Ruas, Meriem Beloucif, Christine de Kock, Nirmal Surange, Daniela Teodorescu, Ibrahim Said Ahmad, David Ifeoluwa Adelani, Alham Fikri Aji, Felermimo D. M. A. Ali, Ilseyar Alimova, Vladimir Araujo, Nikolay Babakov, Naomi Baes, Ana-Maria Bucur, Andiswa Bukula, Guanqun Cao, Rodrigo Tufino Cardenas, Rendi Chevi, Chiamaka Ijeoma Chukwuneke, Alexandra Ciobotaru, Daryna Dementieva, Murja Sani Gadanya, Robert Geislinger, Bela Gipp, Oumaima Hourrane, Oana Ignat, Falalu Ibrahim Lawan, Rooweither Mabuya, Rahmad Mahendra, Vukosi Marivate, Andrew Piper, Alexander Panchenko, Charles Henrique Porto Ferreira, Vitaly Protasov, Samuel Rutunda, Manish Shrivastava, Aura Cristina Udrea, Lilian Diana Awuor Wanzare, Sophie Wu, Florian Valentin Wunderlich, Hanif Muhammad Zhafran, Tianhui Zhang, Yi Zhou, and Saif M. Mohammad. 2025a. [BRIGHTER: BRIdging the Gap in Human-Annotated Textual Emotion Recognition Datasets for 28 Languages](#). Preprint, arXiv:2502.11926.
- Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Seid Muhie Yimam, Jan Philip Wahle, Terry Ruas, Meriem Beloucif, Christine De Kock, Tadesse Destaw Belay, Ibrahim Said Ahmad, Nirmal Surange, Daniela Teodorescu, David Ifeoluwa Adelani, Alham Fikri Aji, Felermimo Ali, Vladimir Araujo, Abinew Ali Ayele, Oana Ignat, Alexander Panchenko, Yi Zhou, and Saif M. Mohammad. 2025b. SemEval task 11: Bridging the gap in text-based emotion detection. In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, Vienna, Austria. Association for Computational Linguistics.
- Esther Rinke and Cristina Flores. 2021. Portuguese as heritage language in Germany—A linguistic perspective. *Languages*, 6(1):10.
- Stephen Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1995. [Okapi at TREC-3](#). In *TREC*.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Tobias Schröder, Kimberly B Rogers, Shuichirou Ike, Julija N Mell, and Wolfgang Scholl. 2013. Affective meanings of stereotyped social groups in cross-cultural comparison. *Group Processes & Intergroup Relations*, 16(6):717–733.
- Hinrich Schutze, Christopher D Manning, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Dorian Stuart. 2022. [Rank BM25](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Elsbeth Turcan, Smaranda Muresan, and Kathleen McKeeown. 2021. Emotion-infused models for explainable psychological stress detection. In *Proceedings of the 2021 conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2895–2909.
- Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, et al. 2024. Aya Model: An instruction finetuned open-access multilingual language model. *arXiv preprint arXiv:2402.07827*.
- Wikipedia. 2025. [Shared heritage of German and English](#). Accessed: 2025-02-26.
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.