AtlasIA at SemEval-2025 Task 11: FastText-Based Emotion Detection in Moroccan Arabic for Low-Resource Settings

Abdeljalil El Majjodi[1], Imane Momayiz[1], and Nouamane Tazi[1]

[1]AtlasIA

Abstract

This study addresses multi-label emotion classification in Moroccan Arabic. We developed a lightweight computational approach to detect and categorize emotional content in seven distinct categories: anger, fear, joy, disgust, sadness, surprise, and neutral. Our findings reveal that our efficient, subword-aware model achieves 46.44% accuracy on the task, demonstrating the viability of lightweight approaches for emotion recognition in under-resourced language variants. The model's performance, while modest, establishes a baseline for emotion detection in Moroccan Arabic, highlighting both the potential and challenges of applying computationally efficient architectures to dialectal Arabic processing. Our analysis reveals particular strengths in handling morphological variations and out-of-vocabulary words, though challenges persist in managing code-switching and subtle emotional distinctions. These results offer valuable insights into the trade-offs between speed and accuracy in multilingual emotion detection systems, particularly for low-resource languages.

1 Introduction

Emotion detection in natural language processing (NLP) presents significant challenges, particularly for low-resource languages like Moroccan Darija (Moroccan Arabic). As digital communication proliferates, understanding emotional nuances becomes crucial for applications in sentiment analysis, social media monitoring, and psychological research (Gandhi et al., 2023). Despite extensive research in emotion classification for well-resourced languages, Moroccan Darija remains under-explored due to several challenges. Its *linguistic complexity*, marked by high context-dependency and significant deviations from Modern Standard Arabic, complicates NLP tasks. The unique *code-switching patterns*, blending Berber, French, Spanish, and Arabic, further hinder traditional approaches (Zaidan and Callison-Burch, 2014). A major bottleneck is the *scarcity of annotated datasets, linguistic tools, and computational resources*. Unlike well-studied languages, Darija lacks sentiment lexicons and annotated corpora necessary for emotion detection. Additionally, *dialectal variability* introduces regional and social variations, making generalization difficult. Lastly, Darija's *nuanced emotional expression*, deeply tied to prosody, idiomatic speech, and cultural context, remains a challenge for conventional NLP techniques.

For these reasons, traditional machine learning approaches often struggle with Darija's unique dialectal characteristics, requiring specialized computational techniques to navigate its intricate linguistic landscape. Developing robust emotion detection models for Darija is not only a technical challenge but also a crucial step toward preserving and understanding the linguistic and emotional nuances of Moroccan Arabic communication.

This paper presents the following contributions:

- **Multi-label Emotion Classification**: We develop[1] a multi-label emotion classification model for Moroccan Darija using FastText (Joulin et al., 2016), achieving **46.44%** accuracy in a low-resource setting.

- **Linguistic Adaptation**: We introduce a specialized preprocessing and modeling approach that handles the linguistic complexities of Moroccan Arabic, including code-switching, non-standard spellings, and dialectal variations.

- **Low-Resource NLP Insights**: We provide insights into the challenges of multi-label emotion detection in under-resourced languages,

[1]Our implementation is open-sourced at https://github.com/atlasia-ma/semeval-emotion-detection

Proceedings of the The 19th International Workshop on Semantic Evaluation (SemEval-2025), pages 34–37
July 31 - August 1, 2025

highlighting the trade-offs between model accuracy and computational constraints.

2 Background

SemEval 2025 Task 11 (Muhammad et al., 2025b) focuses on emotion classification in Moroccan Arabic (Darija), requiring systems to categorize text into seven emotions: *anger, fear, joy, disgust, sadness, surprise*, and *neutral*. For example:

Input: "إليسا: كنعيش قصة حب جديدة ومغاديش نقول لكم معامن"
Label: `joy`

Moroccan Darija presents distinct challenges due to its lack of standardization, extensive dialectal variability, and frequent code-switching with French, Berber, and Spanish. These factors make traditional NLP techniques less effective, requiring specialized preprocessing and modeling strategies. Successfully addressing this task not only advances emotion detection for under-resourced languages but also enhances the accessibility and understanding of sentiment in Moroccan Arabic social media and digital communication.

3 System Overview

Our system leverages FastText (Joulin et al., 2016) for multi-label emotion classification, chosen for its efficiency and ability to handle dialectal variations. A key advantage of FastText is its subword-aware representations, which are particularly beneficial for dialectal Arabic, where words exhibit high morphological variability (Bojanowski et al., 2017). Additionally, its lightweight architecture makes it well-suited for resource-constrained environments, ensuring scalability for real-world applications. Another crucial factor is its robustness to out-of-vocabulary (OOV) words, a common issue in informal Moroccan darija text due to spelling variations, transliterations, and code-switching.

The system follows a structured pipeline, beginning with text preprocessing to clean and normalize input, addressing noise, non-standard spellings, and multilingual elements. The processed text is then used to train the FastText model, leveraging subword embeddings. Finally, during inference, emotion labels are predicted and post-processed to refine outputs, ensuring interpretability and better alignment with the nuances of Moroccan Darija. This approach effectively tackles the linguistic challenges of the task, making emotion detection in Darija more robust and scalable.

4 Data Preparation and Model Training

Our study is based on the SemEval 2025 Task 11 (of SemEval-2025 Task 11, 2025) dataset (Muhammad et al., 2025a), which focuses on emotion detection in Moroccan Arabic (Darija). The dataset is loaded using the Hugging Face `datasets` library, allowing for efficient and reproducible data handling. We specifically work with the Moroccan Arabic ('ary') subset, ensuring that our approach directly addresses the challenges posed by dialectal variation and code-switching in Darija.

4.1 Dataset Statistics

The dataset is relatively small compared to high-resource languages, highlighting the challenges of training robust emotion classification models for underrepresented languages. It consists of:

- **Training set:** 1,608 samples
- **Validation set:** 267 samples
- **Test set:** 812 samples

Each text sample is annotated with one or more emotion labels from seven categories: *anger, fear, joy, disgust, sadness, surprise*, and *neutral*. The multi-label nature of the task reflects the complexity of emotional expression, where a single utterance may convey overlapping sentiments.

4.2 Data Processing Pipeline

To effectively handle noisy and informal text, we employ a structured data processing pipeline using Python. The preprocessing pipeline is designed to normalize noisy user-generated text while preserving relevant linguistic information. The processing in Algorithm 1 was applied.

This approach ensures that irrelevant symbols and formatting inconsistencies are removed, improving the model's ability to generalize across different textual styles.

4.3 Label Processing

Emotion detection in Moroccan Darija is particularly challenging due to the subtle interplay between linguistic and cultural factors. Our multi-label classification framework captures overlapping emotions by assigning one or more labels from the following categories:

Algorithm 1 Tweet Preprocessing

```
Require: tweet: input string containing raw tweet text
Ensure: cleaned and preprocessed tweet string
function PreprocessTweet(tweet)
    // Remove @mentions (user handles)
    tweet ← REGEX_REPLACE(tweet, "@+", "")
    // Remove URLs
    tweet ← REGEX_REPLACE(tweet, "https?:§+|www§+", "")
    // Remove hashtags, including hashtagged words
    tweet ← REGEX_REPLACE(tweet, "+", "")
    // Normalize whitespace
    tweet ← REGEX_REPLACE(tweet, "+", " ")
    tweet ← TRIM(tweet)
    // Normalize punctuation (remove excessive spaces around punctuation)
    tweet ← REGEX_REPLACE(tweet, "*([.,!?;:])*", "  ")
return tweet
end function
```

- **Primary emotions:** anger, fear, joy, disgust, sadness, surprise
- **Neutral category:** for texts that do not explicitly express emotion

Given the inherent subjectivity of emotion annotation, the model must learn to distinguish between subtle variations in sentiment while accounting for dialectal expressions, code-switching, and informal speech patterns. By structuring our label processing accordingly, we enhance the model's ability to capture the nuances of Moroccan Arabic emotional expression.

4.4 Model Configuration

The FastText model was trained with the optimized parameters outlined in Table 1.

5 Results

5.1 Model Performance

Our system achieved a validation accuracy of **46.44%**, demonstrating the viability of using FastText for emotion classification in Moroccan Arabic. Despite the challenges of working with a low-resource language, this result serves as a promising baseline for future research and model improvements in the context of under-resourced dialectal languages.

Parameter	Value
Learning rate	0.40
Dimension	8
Window size	5
Epochs	100
Min word count	1
Character n-grams	3-6
Word n-grams	1
Loss function	One-vs-All (OVA)
Bucket size	2,247,558
Threads	7
Learning rate update rate	100
T value	0.0001

Table 1: FastText model configuration parameters

5.2 Analysis of Errors

A thorough qualitative error analysis revealed several factors contributing to model misclassifications:

- **Handling of Code-Switched Text:** The model struggles with tweets containing multiple languages, particularly when sentiment-laden words are in French or Berber rather than Arabic.
- **Ambiguity Between Similar Emotions:** Emotions such as *anger* and *surprise*, or *fear* and *sadness*, share overlapping linguistic markers and contextual cues, making differentiation challenging.
- **Performance on Multi-Label Cases:** The model tends to focus on the most dominant emotion, neglecting secondary emotions when multiple labels are present in a sample.

For example, the following error case illustrates the model's difficulty in distinguishing between closely related emotions:

Input:
"بنكي فكازا ضرب فلوس صحيحة لكايان"
Predicted: anger
True label: surprise
Analysis: The model misclassified *surprise* as *anger*, likely due to the intense tone and overlapping intensity between

the two high-arousal emotions. This example underscores the challenges in capturing subtle emotional distinctions in informal, dialectal speech.

These insights point to opportunities for improving the model's handling of ambiguous emotions and code-switching, suggesting potential avenues for enhancing the model's robustness and accuracy in future iterations.

6 Conclusion and Future Work

Our emotion detection system for Moroccan Arabic faces several limitations, including difficulty handling code-switching between Arabic, French, and Berber, reduced performance on informal text variations, and challenges in capturing context-dependent emotional nuances. The current model architecture also struggles with Darija's complex morphology, and the evaluation is limited to accuracy, which does not fully reflect the model's performance.

Future work should focus on developing emotion lexicons specific to Moroccan Arabic, integrating Darija-specific linguistic features, and exploring hybrid approaches combining FastText with transformer models like AraBERT (Antoun et al., 2020). Cross-lingual transfer learning, larger and more diverse emotion datasets, and the adoption of advanced evaluation metrics will further improve the model's effectiveness and robustness.

Acknowledgments

We thank the SemEval 2025 Task 11 (of SemEval-2025 Task 11, 2025) organizers for providing the dataset and evaluation framework. We also acknowledge the Hugging Face team for their datasets library and the FastText(Joulin et al., 2016) team for their efficient implementation.

References

Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. AraBERT: Transformer-based model for Arabic language understanding. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15, Marseille, France. European Language Resource Association.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Preprint*, arXiv:1607.04606.

Ankita Gandhi, Kinjal Adhvaryu, Soujanya Poria, Erik Cambria, and Amir Hussain. 2023. Multimodal sentiment analysis: A systematic review of history, datasets, multimodal fusion methods, applications, challenges and future directions. *Information Fusion*, 91:424–444.

Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.

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, Felermino 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, Felermino 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.

Organizers of SemEval-2025 Task 11. 2025. Semeval-2025 task 11: Bridging the gap in text-based emotion detection (track a). In *Proceedings of the 14th International Workshop on Semantic Evaluation (SemEval-2025)*. Association for Computational Linguistics.

Omar F. Zaidan and Chris Callison-Burch. 2014. Arabic dialect identification. *Computational Linguistics*, 40(1):171–202.