

Deep Multi-Task Model for Sarcasm Detection and Sentiment Analysis in Arabic Language

Abdelkader El Mahdaouy¹ Abdellah El Mekki¹ Nabil El Mamoun²
Kabil Essefar¹ Ismail Berrada¹ Ahmed Khoumsi³

¹School of Computer Sciences, Mohammed VI Polytechnic University, Morocco

²Faculty of Sciences Dhar EL Mahraz, Sidi Mohamed Ben Abdellah University, Morocco

³Dept. Electrical & Computer Engineering, University of Sherbrooke, Canada

{firstname.lastname}@um6p.ma

ahmed.khoumsi@usherbrooke.ca

Abstract

The prominence of figurative language devices, such as sarcasm and irony, poses serious challenges for Arabic Sentiment Analysis (SA). While previous research works tackle SA and sarcasm detection separately, this paper introduces an end-to-end deep Multi-Task Learning (MTL) model, allowing knowledge interaction between the two tasks. Our MTL model's architecture consists of a Bidirectional Encoder Representation from Transformers (BERT) model, a multi-task attention interaction module, and two task classifiers. The overall obtained results show that our proposed model outperforms its single-task counterparts on both SA and sarcasm detection sub-tasks.

1 Introduction

The popularity of the Internet and the unprecedented reach of social media platforms allow users to express their opinions on a wide range of topics. Thereby, Sentiment Analysis (SA) has become a cornerstone for many applications such as digital marketing, product review analysis, customer feedback, social media monitoring, etc. SA consists of determining the expressed sentiment (positive, negative, or neutral) conveyed by a text or a piece of text.

Over the past decade, significant research advances have been achieved for Arabic SA (Badaro et al., 2019; Al-Ayyoub et al., 2019; Oueslati et al., 2020; Abu Farha and Magdy, 2021). However, the mutual interaction and impact of figurative language devices, like sarcasm and irony, and Arabic SA remain under-explored (Abu Farha and Magdy, 2020, 2021; Abbes et al., 2020). These latter devices allow us to express ourselves intelligently beyond the literal meaning of words. Although the literature uses the terms irony and sarcasm interchangeably, they have different meanings and there

is no consensus on their definition (Farías et al., 2016; Hernández Farías and Rosso, 2017; Zhang et al., 2019). Both sarcasm and irony devices pose a real challenge for SA as they can reverse the expressed sentiment polarity from positive to negative (Hernández Farías and Rosso, 2017; Abu Farha and Magdy, 2020, 2021). Therefore, there is an urgent need to develop sarcasm-aware SA tools.

Previous research works on Arabic SA and sarcasm detection have dealt with both tasks separately (Ghanem et al., 2019, 2020; Abbes et al., 2020; Abu Farha and Magdy, 2020). Abbes et al. (2020) have built a corpus for irony and sarcasm detection in Arabic language from twitter using a set of ironic hashtags. Unlike the previous work of Karoui et al. (2017) that have relied on ironic hashtags to label the tweets, the annotation is performed manually by two Arabic language specialists. Abu Farha and Magdy (2021) have presented an overview of existing Arabic SA methods and approaches, and a benchmarking using three existing datasets. Their results have shown that most of the evaluated models perform poorly on the SemEval and ASTD datasets. Due to the label inconsistencies discovered, they have re-annotated the previously mentioned datasets for SA and sarcasm detection. In addition to the highly subjective nature of SA task, they have reported a large performance drop in the case of sarcastic tweets (Abu Farha and Magdy, 2020, 2021).

Following the recent breakthroughs in Arabic Natural Language Processing (NLP), achieved using AraBERT model (Antoun et al., 2020), Abdul-Mageed et al. (2020) have introduced two Arabic transformer-based language models, namely ARBERT and MARBERT. ARBERT is trained on large textual corpora of Modern Standard Arabic (MSA), while MARBERT is trained on 1 billion DA and MSA tweets corpus. They have shown new cutting edge performances on wide range of DA

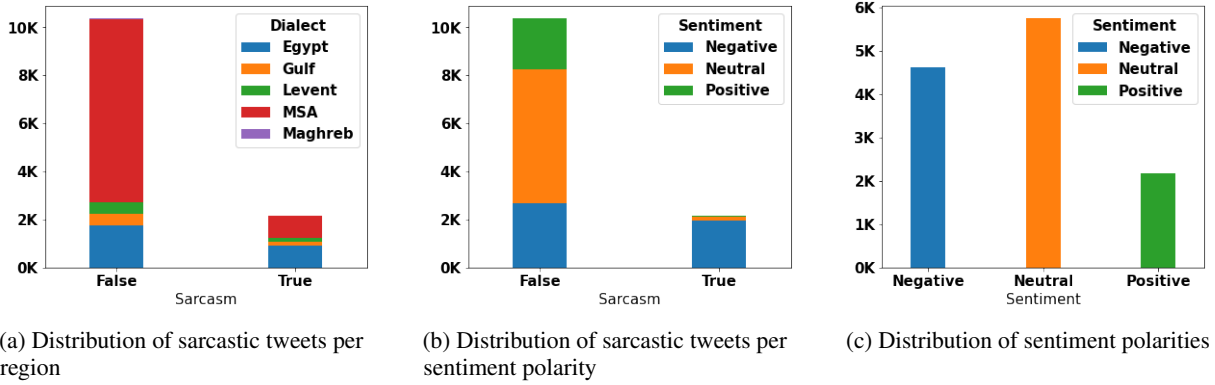


Figure 1: ArSarcasm-v2 dataset: distribution of sarcastic tweets and their sentiment polarities. True and False denote sarcastic and non-sarcastic tweets, respectively.

and MSA NLP tasks (AraBench datasets), including, among others, SA and sarcasm detection.

In this paper, we present our end-to-end deep MTL model, submitted to SA and sarcasm detection in Arabic language shared task (Abu Farha et al., 2021). Our approach is based on MARBERT (Abdul-Mageed et al., 2020), and a multi-task attention interaction module. The latter consists of two task-specific attention layers for extracting task-discriminative features, and of a Sigmoid interaction layer (Lan et al., 2017) for allowing interaction and knowledge sharing between sarcasm detection and SA. The task-interaction is performed using the task-specific attention outputs, a learnable shared matrix, and the Sigmoid activation. The obtained results show that our MTL model surpasses the other evaluated single-task and MTL models. Besides, the incorporation of an attention mechanism and the task-interaction boosts the performance of both sarcasm detection and SA.

The rest of the paper is organized as follows. Section 2 presents the shared task’s dataset. Section 3 introduces the proposed method. In Section 4, we present the obtained results for both sarcasm detection and SA subtasks. Section 5 discusses the obtained results. Finally, Section 6 concludes the paper.

2 Data

The ArSarcasm Shared Task consists of two subtasks for sarcasm detection and SA in Arabic language (Abu Farha et al., 2021). The shared task’s dataset, ArSarcasm-v2, is built from the previously introduced datasets for sarcasm and irony detection (Abbes et al., 2020; Abu Farha and Magdy, 2020). The provided dataset consists of 12,548 and 3,000 tweets for the training set and test set, respec-

tively. The task’s dataset is annotated for SA and sarcasm detection as well as the regional dialect of the tweets.

Figure 1 presents the distribution of sarcastic tweets and their sentiment polarities (Figures 1a and 1c). The distribution of all sentiment polarities in the dataset is illustrated in 1c. The dataset is unbalanced for both subtasks. Most sarcastic tweets are written in MSA and Egyptian dialect (Figure 1a), and are labeled with a negative sentiment (Figure 1b). Furthermore, approximately half of the tweets convey a neutral sentiment (Figure 1c).

3 Method

Our multi-task model consists of three main components: BERT encoder, a multi-task attention interaction module, and two task classifiers.

3.1 BERT Encoder

Fine-tuning Bidirectional Encoder Representation from Transformers (BERT) model on downstream tasks has shown a new wave of state-of-the-art performances in many NLP applications (Devlin et al., 2019). BERT model’s architecture consists of multiple transformer encoders for learning contextualized word embedding of a given input text. It is trained on large textual corpora using two self-supervised objectives, namely the Masked Language Model (MLM) and the Next Sentence Prediction (NSP).

The encoder of our MTL model is the pre-trained MARBERT (Abdul-Mageed et al., 2020). MARBERT is fed with a sequence of wordpieces $[t_1, t_2, \dots, t_n]$ of the input tweet, where n is the sequence length. It outputs the tweet embedding $h_{[CLS]}$ ([CLS] token embedding) and the contextualized word embedding of the input tokens

$H = [h_1, h_2, \dots, h_n] \in \mathbb{R}^{n \times d}$. Both $h_{[CLS]}$ and h_i have the same hidden dimension d .

3.2 Multi-task attention interaction module

This module consists of two task-specific attention layers (task-specific context-rich representation) and a Sigmoid task-interaction layer.

The task-specific sentence representation $v_* \in \mathbb{R}^{1 \times d}$ (e.g. v_{sarc} and v_{sent}) is obtained using the attention mechanism over the contextualized word embedding matrix H :

$$C = \tanh(HW^a)$$

$$\alpha = \text{softmax}(C^T W^\alpha)$$

$$v_* = \alpha \cdot H^T$$

where $W^a \in \mathbb{R}^{d \times 1}$ and $W^\alpha \in \mathbb{R}^{n \times n}$ are the learnable parameters of the attention mechanism. $C \in \mathbb{R}^{n \times 1}$ and $\alpha \in [0, 1]^n$ weights words hidden representations according to their relevance to the task.

The task interaction mechanism (Lan et al., 2017) is performed using a learnable shared matrix $W^i \in \mathbb{R}^{d \times d}$ and a bias vector $b^i \in \mathbb{R}^d$. The interaction of both task are given by:

$$v'_{sarc} = v_{sarc} \odot \sigma(W^i v_{sent} + b^i) \quad (1)$$

$$v'_{sent} = v_{sent} \odot \sigma(W^i v_{sarc} + b^i) \quad (2)$$

where v_{sarc} and v_{sent} are the output of the sarcasm task-specific attention layer and the sentiment task-specific attention layer, respectively. \odot is the element-wise product.

3.3 Task classifier

We employ two task classifiers F_{sarc} and F_{sent} for sarcasm detection and SA, respectively. Each classifier consists of one hidden layer and one output layer. They are fed with the concatenation of the pooled output embedding and the task output of the multi-task attention interaction module v'_* (e.g. v'_{sarc} and v'_{sent}). The outputs of the task classifiers are given by:

$$\hat{y}_{sarc} = F_{sarc}([h_{[CLS]}, v'_{sarc}]) \quad (3)$$

$$\hat{y}_{sent} = F_{sent}([h_{[CLS]}, v'_{sent}]) \quad (4)$$

3.4 Multi-task learning objective

We train our MTL model to jointly minimize the binary cross-entropy loss \mathcal{L}_{BCE} , for sarcasm detection, and the cross-entropy loss \mathcal{L}_{CE} , for SA. The total loss is given by:

$$\mathcal{L} = \mathcal{L}_{BCE}(y_{sarc}, \hat{y}_{sarc}) + \mathcal{L}_{CE}(y_{sent}, \hat{y}_{sent}) \quad (5)$$

where \hat{y}_* is the predicted output and y_* is the ground truth label.

4 Results

In this section, we present the experiment settings and the obtained results.

4.1 Experiment settings

We have compared our model (MTL_ATTINTER) with two single-task models (ST and ST_ATT) and two MTL models (MTL and MTL_ATT).

- **ST** consists of MARBERT with one classification layer.
- **ST_ATT** employs the attention mechanism on top of the contextualized word embedding of MARBERT. The classification is performed using the attention layer output and the [CLS] token embedding.
- **MTL** is similar to **ST** model and uses classification layer for each task.
- **MTL_ATT** is the MTL counterpart of **ST_ATT** model.

We have implemented the MARBERT' tweets preprocessing pipeline (Abdul-Mageed et al., 2020). The evaluated models have been trained using Adam optimizer with a learning rate of 5×10^{-6} . Based on several experiments, the batch size and the number of epochs have been fixed to 64 and 5, respectively. Besides, we have used 80% and 20% of the provided training data for training set and development set, respectively. For comparison purposes, we have used the macro-average Precision, Recall, F1, and F1 score of positive and negative ($F1^{PN}$) evaluation measures. We have also employed the Accuracy and the F1 score of the sarcastic tweets ($F1^{Sarc}$).

4.2 Experiment results

Table 1 shows the obtained models' performances for both SA and sarcasm detection. The best results, for each evaluation measure, are highlighted with

		Sarcasm					Sentiment				
		Precision	Recall	Accuracy	F1	F1 ^{Sarc}	Precision	Recall	Accuracy	F1	F1 ^{PN}
ST	Dev	0.7649	0.7683	0.8673	0.7666	0.6132	0.7422	0.7519	0.7641	0.7465	0.7284
	Test	0.706	0.708	0.768	0.707	0.573	0.672	0.667	0.713	0.665	0.749
ST_ATT	Dev	0.7736	0.7588	0.8622	0.7658	0.6156	0.7541	0.7429	0.7629	0.7479	0.7253
	Test	0.724	0.722	0.778	0.723	0.598	0.664	0.665	0.709	0.661	0.742
MTL	Dev	0.7935	0.7611	0.8633	0.7753	0.6347	0.7424	0.748	0.7649	0.7448	0.7288
	Test	0.725	0.714	0.771	0.719	0.599	0.676	0.656	0.703	0.662	0.736
MTL_ATT	Dev	0.8064	0.7581	0.8606	0.7778	0.6421	0.7478	0.7524	0.7649	0.7465	0.7326
	Test	0.741	0.72	0.773	0.728	0.617	0.663	0.676	0.717	0.66	0.752
MTL_ATTINTER	Dev	0.8106	0.766	0.8661	0.7846	0.6522	0.7511	0.7414	0.7582	0.7436	0.7358
	Test	0.7268	0.7122	0.7680	0.7183	0.6000	0.6713	0.7183	0.7107	0.6625	0.7480

Table 1: Models evaluation on both SA and sarcasm detection subtasks

		Sarcasm					Sentiment				
		Precision	Recall	Accuracy	F1	F1 ^{Sarc}	Precision	Recall	Accuracy	F1	F1 ^{PN}
MTL_ATTINTER		0.7268	0.7122	0.7680	0.7183	0.6000	0.6713	0.7183	0.7107	0.6625	0.7480

Table 2: The obtained results of our **Official** submission

italic font and bold font for the dev set and the test set, respectively. The overall obtained results show that MTL models outperform their single-task counterparts for most evaluation measures. In fact, incorporating attention mechanism into both ST_ATT and MTL_ATT improves the F1, F1^{Sarc} and F1^{PN}. The former compute the F1 score for sarcastic tweets only, while the latter consider only positive and negative sentiment.

MTL_ATTINTER and MTL_ATT achieve the best performances for most evaluation measures on both the dev and the test sets of sarcasm detection sub-task. Specifically, they show far better F1 performance for the sarcastic class prediction. For SA, the other evaluated models achieve slightly better performance. However, MTL_ATTINTER and MTL_ATT yield the best F1^{PN} performances on the dev set and the test set. Therefore, our proposed model excels in detecting sarcastic tweets as well as predicting positive and negative sentiments.

Official results

Since one submission was allowed, we have submitted the results of our MTL_ATTINTER model. Table 2 shows the official submission results. Our system is top ranked on SA Sub-task and has secured the fourth position among submitted systems for sarcasm detection.

5 Discussion

To investigate the strengths and weaknesses of our model, we have analyzed the confusion matrix of each subtask (Figures 2a and 2b) as well as the confusion matrices of sentiment analysis among sarcastic and non-sarcastic tweets respectively (Figures 2d and 2c). The analysis of these matrices shows that our MTL model leverages signals from both tasks and boosts the performances. This can be explained by the fact that most sarcastic tweets convey a negative sentiment. Besides, negative tweets tend to have a large probability of being sarcastic than the positive ones. This could be also deduced from Table 1, where MTL models achieve the best F1^{Sarc} and F1^{PN} scores compared to single-task models.

6 Conclusion

In this paper, we have proposed an end-to-end deep Multi-Task Learning model for SA and sarcasm detection. Our model leverages the MARBERT’s contextualized word embedding with a multi-task attention interaction module. The aim is to allow task-interaction and knowledge sharing for both SA and sarcasm detection. Our model shows very promising results on both subtasks. Therefore, it proves the effectiveness of using task-specific attention layers as well as the task-interaction mechanism in multi-task learning.

Future research work will focus on developing

False	1859	217
True	119	315
	False	True

(a) Confusion matrix of the sarcasm detection task

Negative	787	98	40
Neutral	222	813	114
Positive	51	82	303
	Negative	Neutral	Positive

(b) Confusion matrix of SA task

Negative	430	76	31
Neutral	196	807	112
Positive	45	81	298
	Negative	Neutral	Positive

(c) Confusion matrix of SA among non-sarcastic tweets

Negative	357	22	9
Neutral	26	6	2
Positive	6	1	5
	Negative	Neutral	Positive

(d) Confusion matrix of SA among sarcastic tweets

Figure 2: The confusion matrices of our MTL model’s prediction on both SA and sarcasm detection tasks

task-interaction and class-interaction modules and mechanisms for SA and sarcasm detection.

References

- Ines Abbes, Wajdi Zaghouni, Oaima El-Hardlo, and Faten Ashour. 2020. [DAICT: A dialectal arabic irony corpus extracted from twitter](#). In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, pages 6265–6271. European Language Resources Association.
- Muhammad Abdul-Mageed, AbdelRahim Elmadany, and El Moatez Billah Nagoudi. 2020. [Arbert & marbert: Deep bidirectional transformers for arabic](#). *arXiv preprint arXiv:2101.01785*.
- Ibrahim Abu Farha and Walid Magdy. 2020. [From arabic sentiment analysis to sarcasm detection: The arsarcasm dataset](#). In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 32–39.
- Ibrahim Abu Farha and Walid Magdy. 2021. [A comparative study of effective approaches for arabic sentiment analysis](#). *Information Processing Management*, 58(2):102438.
- Ibrahim Abu Farha, Wajdi Zaghouni, and Walid Magdy. 2021. [Overview of the wanlp 2021 shared task on sarcasm and sentiment detection in arabic](#). In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*.
- Mahmoud Al-Ayyoub, Abed Allah Khamaiseh, Yaser Jararweh, and Mohammed N. Al-Kabi. 2019. [A comprehensive survey of arabic sentiment analysis](#). *Information Processing & Management*, 56(2):320–342. Advance Arabic Natural Language Processing (ANLP) and its Applications.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. [Arbert: Transformer-based model for arabic language understanding](#). In *LREC 2020 Workshop Language Resources and Evaluation Conference 11–16 May 2020*, page 9.
- Gilbert Badaro, Ramy Baly, Hazem Hajj, Wassim El-Hajj, Khaled Bashir Shaban, Nizar Habash, Ahmad

- Al-Sallab, and Ali Hamdi. 2019. [A survey of opinion mining in arabic: A comprehensive system perspective covering challenges and advances in tools, resources, models, applications, and visualizations.](#) *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 18(3).
- 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.
- Delia Irazú Hernández Farías, Viviana Patti, and Paolo Rosso. 2016. [Irony detection in twitter: The role of affective content.](#) *ACM Trans. Internet Technol.*, 16(3).
- Bilal Ghanem, Jihen Karoui, Farah Benamara, Véronique Moriceau, and Paolo Rosso. 2019. [Idat at fire2019: Overview of the track on irony detection in arabic tweets.](#) In *Proceedings of the 11th Forum for Information Retrieval Evaluation, FIRE '19*, page 10–13, New York, NY, USA. Association for Computing Machinery.
- Bilal Ghanem, Jihen Karoui, Farah Benamara, Paolo Rosso, and Véronique Moriceau. 2020. [Irony detection in a multilingual context.](#) In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part II*, volume 12036 of *Lecture Notes in Computer Science*, pages 141–149. Springer.
- Delia Irazú Farias Hernández Farías and Paolo Rosso. 2017. [Chapter 7 - irony, sarcasm, and sentiment analysis.](#) In Federico Alberto Pozzi, Elisabetta Fersini, Enza Messina, and Bing Liu, editors, *Sentiment Analysis in Social Networks*, pages 113–128. Morgan Kaufmann, Boston.
- Jihen Karoui, Farah Banamara Zitoune, and Véronique Moriceau. 2017. [Soukhria: Towards an irony detection system for arabic in social media.](#) *Procedia Computer Science*, 117:161–168. Arabic Computational Linguistics.
- Man Lan, Jianxiang Wang, Yuanbin Wu, Zheng-Yu Niu, and Haifeng Wang. 2017. [Multi-task attention-based neural networks for implicit discourse relationship representation and identification.](#) In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1299–1308. Association for Computational Linguistics.
- Oumaima Oueslati, Erik Cambria, Moez Ben HajHmida, and Habib Ounelli. 2020. [A review of sentiment analysis research in arabic language.](#) *Future Generation Computer Systems*, 112:408 – 430.
- Shiwei Zhang, Xiuzhen Zhang, Jeffrey Chan, and Paolo Rosso. 2019. [Irony detection via sentiment-based transfer learning.](#) *Information Processing Management*, 56(5):1633–1644.