
ASAD: A TWITTER-BASED BENCHMARK ARABIC SENTIMENT ANALYSIS DATASET

Basma Alharbi*
Univ. of Jeddah
Saudi Arabia

Hind Alamro*
KAUST
Saudi Arabia

Manal Alshehri*
KAUST
Saudi Arabia

Zuhair Khayyat
Lucidya
Saudi Arabia

Manal Kalkatawi
KAU
Saudi Arabia

Inji Ibrahim Jaber
KAUST
Saudi Arabia

Xiangliang Zhang^{✉†}
KAUST
Saudi Arabia

ABSTRACT

This paper provides a detailed description of a new Twitter-based benchmark dataset for Arabic Sentiment Analysis (ASAD), which is launched in a competition³, sponsored by KAUST for awarding 10000 USD, 5000 USD and 2000 USD to the first, second and third place winners, respectively. Compared to other publicly released Arabic datasets, ASAD is a large, high-quality annotated dataset (including 95K tweets), with three-class sentiment labels (*positive*, *negative* and *neutral*). We presents the details of the data collection process and annotation process. In addition, we implement several baseline models for the competition task and report the results as a reference for the participants to the competition.

Keywords Sentiment Analysis · Arabic Tweets · Deep Learning · Benchmark Dataset

1 Introduction

Sentiment analysis is a popular and challenging Natural Language Processing (NLP) task, which facilitates capturing public opinions on a topic, product or a service. It is thus a widely studied research area, where many benchmark datasets and algorithms exist to tackle this challenge. The state of research in Arabic sentiment analysis falls behind other high-source languages such as English and Chinese [1]. A comprehensive review on the current state-of-research on Arabic sentiment analysis can be found in [2]. In general, the work on Arabic sentiment analysis can be grouped by the employed approaches, which range from simple, rule-based / lexicon-based algorithms [3] to more complex deep learning algorithms [4]. Synthesizing recent publications on Arabic sentiment analysis show that there is an increase in the utilization of deep learning approaches in the past few years [5, 6, 7, 8, 9, 10]. In such approaches, often a combination of word embedding model for feature representation is used along with a deep learning model for classification. For example, in [10], the authors applied three deep learning models: CNN, LSTM and RCNN, and the input to all these models were word embeddings generated from a pretrained word2vec CBOW model [11], named Aravec [12]. The contributions of work in Arabic deep learning-based sentiment analysis is limited to the application of existing models to Arabic corpora, and only a few publications contribute new models that are specifically designed to target the unique characteristics of the Arabic language, as in [13].

The lack of large and high-quality Arabic benchmark dataset is often attributed to this deficiency in Arabic-based NLP studies [14]. The objective of this work is to introduce a new large Twitter-based Arabic Sentiment Analysis Dataset

*The first three authors had equal contributions.

^{†✉}The corresponding author: Xiangliang Zhang, Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division and Computational Biosciences Research Center (CBRC), King Abdullah University of Science and Technology (KAUST). Email: xiangliang.zhang@kaust.edu.sa

³<https://wti.kaust.edu.sa/arabic-sentiment-analysis-challenge>

Table 1: Summary of Twitter-based Arabic Sentiment Analysis Corpora. The '-' sign indicates that these metrics were not explicitly specified in the paper.

	Size	No. Classes	Dialect	Annotation Approach	No. of Annotators	Annotation Quality
TEAD [20]	5,500 M	3	Multi	from Emojis	-	-
40K Tweets [10]	40,000	2	MSA, Egyptian	-	-	-
ASTAD [21]	36,000	2	Multi	from Emojis	-	-
Ara-SenTi-Tweet [22]	17,573	4	MSA, Saudi	Manual	3	$\kappa = 0.60$
ASTD [25]	10,000	4	Egyptian	Manual	3	-
Gold Standard [23]	8,868	4	Multi	Manual	2	$\kappa = 0.82$
SemEval [24]	9,455	5	NA	Manual	5	-
Our ASAD	100,000	3	Multi	Manual	3	$\kappa = 0.56$

(ASAD). ASAD is a public dataset intended to accelerate research in Arabic NLP in general, and Arabic sentiment classification in specific. A competition⁴ is launched on top of the ASAD dataset.

This being said, the increasing interest in the past few years in Arabic sentiment analysis has resulted in the release of a number of datasets. Datasets collected for Arabic sentiment analysis span multiple sources, which include newspapers, reviews, and different content from social media platforms [15, 16, 17, 18]. Table 1 compares our dataset ASAD with similar benchmark datasets, which are also Twitter-based Arabic corpora for general sentiments (as opposed to topic-based sentiment). The benchmark datasets are compared in terms of the size of dataset (i.e., the number of annotated tweets), the number of sentiment classes, the inclusion of Arabic dialect, the annotation approach, the number of annotators and the quality of the annotation (measured by Fleiss’s Kappa metric [19]).

TEAD [20], the currently largest dataset for Arabic sentiment analysis, was annotated automatically using emojis in tweets. The same automatic annotation approach was adopted by ASTAD [21], while no explicit details of how the 40k Tweets dataset was annotated [10]. The remaining datasets have significantly smaller size, yet they were all annotated manually. Manual annotations indicate that a number of human annotators were recruited and trained to provide high quality annotations for each record in the dataset. Manual annotation thus provides a more reliable data source for sentiment classification algorithms to learn from. The quality of annotations can be assessed, e.g., by the most commonly used Fleiss’s Kappa metric [19], whose values closer to 1 indicate higher agreements between annotators. Among the four manually annotated corpora, AraSenTi-Tweet [22] and Gold Standard [23] stated their annotation quality measure, which are included in Table 1.

SemEval [24] annotates tweets in either one of five sentiments, which are: *strongly positive*, *weakly positive*, *neutral*, *weakly negative* and *strongly negative*. AraSenTi-Tweet [22], ASTD [25], and Gold standard [23] adopt a four-point scale for sentiments, which are: *positive*, *neutral*, *negative* and *mixed classes*. Lastly, TEAD [20] follows a three-point scale for sentiments, where each tweet is either *positive*, *negative* or *neutral*. Both 40K Tweets [10] and ASTAD [21] use a two-point scale for sentiment, where tweets are either positive or negative. Our dataset, ASAD, adopts a three-point scale similarly to TEAD [20], has multiple dialects, and was annotated manually by an average of three annotators per tweet. The annotation quality for our dataset is $\kappa = 0.56$, indicating a moderate agreement between the annotators. A detailed analysis of the annotation is presented in the next section.

The rest of this paper is organized as follows. Section 2 provides details on the construction mechanism we followed to collect ASAD. Section 3 provides details on benchmark evaluation and results using baseline approaches for three-class sentiment classification. Section 4 provides a discussion on challenges and future applications to ASAD.

2 Corpus Construction

Figure 1 illustrates the methodology we adopted to construct the corpus. In general, the methodology consists of three main phases: data collection, data annotation and data cleaning. Each phase is further described in the subsections below. The data collection and annotation was done by Lucidya⁵ which is an AI-based company with rich experience in organizing data annotation projects.

⁴Competition website: <https://wti.kaust.edu.sa/arabic-sentiment-analysis-challenge>

⁵<https://lucidya.com/>

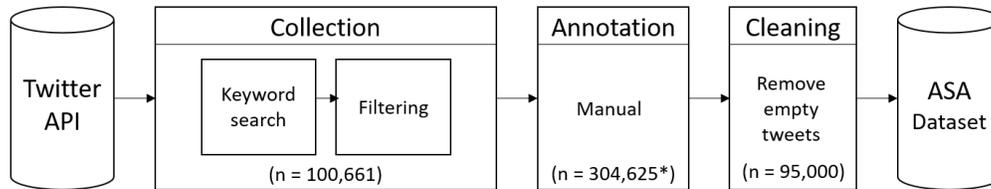


Figure 1: Corpus Collection Methodology. Illustration of the different phases involved in collecting, annotating and cleaning the dataset. * In the annotation phase, each record was annotated by at least three annotators. This justifies the increase in the number of records in the annotation phase when compare to the previous phase (the data collection phase).

2.1 Corpus Collection

The tweets in ASAD were randomly selected from a pool of tweets that have been collected using the Twitter public streaming API⁶ between May 2012 and April 2020. The selection process consisted of several filtration operations to ensure that the Arabic texts are comprehensive enough for the annotation task. First, a random set of tweets with "ar" label in their json object were selected from the pool of tweets and added to the corpus. Next, tweets that contained less than seven words (excluding hashtags and user mentions) were removed from the corpus. After that, all tweets that included URLs, images, or videos or consisted of a predefined set of inappropriate Arabic keywords were also removed from the Corpus. Finally, the corpus was refined to include 100k tweets.

We analyzed the content of ASAD and found out that around 69% of the tweets were selected from the year 2020 (January 2020 until April 2020), 30% were tweeted in 2019 and the remaining 1% happened between 2012 and 2018. 82% of the tweets came from unique authors where we encountered only 503 verified authors in the corpus. For locations, 72% of the tweets originated from Saudi Arabia, 13% from Egypt, 7% from Kuwait, 3% from the United States, and 3% from the United Arab Emirates. Based on the models developed by Lucidya, 36% of the tweets are written in Modern Arabic, 31% of the tweets used the Khaleeji dialect, 22% used the Hijazi dialect, and the rest 10% tweeted in Egyptian dialect.

2.2 Corpus Annotation

The objective of this phase is to apply tweet-level sentiment annotation. We adopted a three-way classification sentiment, where a sentiment can either be: *positive*, *negative* or *neutral*. A screenshot of the platform for our project is given in Figure 2. Each tweet was annotated by at least three experienced annotators who were trained and qualified by the company.

The annotation process was done by 69 native Arabic speakers with an average age of 26 years old where 12% of the recruited annotators were under the age of 21, 71% between 21 and 30, 13% between 31 and 40, and the rest over 40 years old. Among the annotators, 85% of them were females, and 71% of them have at least a Bachelor's degree, and 4% obtained a degree in higher education. In addition, 67% of the annotators have a background in arts, linguistics, or education while the remaining 33% majored in either science or engineering. Based on Lucidya's experience, it is more efficient to hire local annotators which resulted in having 71% of the annotators to be from Saudi Arabia. The remaining annotators came from various parts of the Middle East where 12% were hired from Egypt, 7% from Yemen, 7% from Algeria, while the rest came from Indonesian and Palestine.

The annotators started the corpus annotation project on February 2020 and finished by May 2020 where the tweets were added to the the annotation portal in several batches⁷. By the end of each batch, an evaluation of the quality of the annotation was conducted to ensure that only high quality annotators were allowed to participate in the next batch. By the end of the 10 batches, we gained a total of 304,625 annotations for 100,661 distinct tweets, with a ratio of 3 annotations per tweet.

2.3 Data Cleaning

Preliminary data cleaning was done to ensure minimum data quality level. Empty tweets, i.e., tweets with no text, as well as meaningless tweets (e.g., those containing only URLs or coupons) were excluded from the corpus. The total number of identified records at this stage is 95,000.

⁶<https://developer.twitter.com/en/docs/twitter-api>

⁷Each batch of size 10,000 tweets.

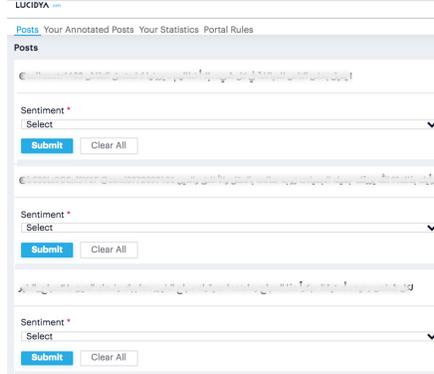


Figure 2: Screenshot of the Annotation Platform.

Table 2: The Fleiss Kappa coefficient κ of ASAD

No. annotators per tweet	No. tweets (N)			κ
	with 1 agreed label	with 2 different labels	with 3 different labels	
$n=3$	58424 (62.34%)	35285 (37.65%)	10 (0.01%)	0.5035
$n=4$	813 (63.47%)	439(34.27%)	29(2.26%)	0.6067
averaged				$\kappa=0.56$

2.3.1 Annotation Evaluation

For a corpus to be considered as a benchmark, we need to verify the reliability of the labels as well as the difficulty level of the task. For that, we adopt the Fleiss’s Kappa [19] metric to measure the Inter Annotator Agreement (IAA). The Fleiss’s Kappa measure is calculated by

$$\kappa = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}; \text{ where} \quad (1)$$

$$\bar{P} = \frac{1}{N} \sum_{i=1}^N \frac{1}{n(n-1)} \sum_{j=1}^k n_{ij}(n_{ij} - 1) \quad (2)$$

$$\bar{P}_e = \sum_{j=1}^k \left(\frac{1}{Nn} \sum_{i=1}^N n_{ij} \right)^2 \quad (3)$$

where N is the total number of tweets, n is the number of annotations per tweet, k is the number of classes, and n_{ij} is the number of annotators who have assigned the i -th tweet to the j -th class. In our ASAD dataset, we have tweets annotated by three annotators ($n=3$), and by four annotators ($n=4$). The agreement among the annotators and the calculated κ for each case is shown in Table 2. The over all Fleiss Kappa coefficient of ASAD is $\kappa=0.56$, indicating a moderate agreement among the annotators.

Since there are around 37% tweets initially annotated with more than two different labels (as given in Table 2), we investigate whether the disagreement is due to the fact that the labeling task is non-trivial for human annotators, or is caused by the carelessness of certain annotators. We determine the final label for each tweet by majority voting or another experienced annotator if there was no consensus. Then, we evaluate the accuracy of each annotator: how much percentage of their annotated tweets are the same as the final label. The accuracy is shown in Figure 3, where the 69 annotators are ordered by their number of contributed annotations. We can see that the top-10 annotators contributed more than 60% annotations with an average accuracy 0.88. The overall average accuracy is 0.85 regardless the number of contributed annotations. Only a few annotators had accuracy less than 0.8, and most of these annotators gave less than 200 annotations. Following [26] and [27], we also evaluate the average inter-rater agreement $\iota=0.83$, indicating the reliability of the annotators since ι is much higher than 0.5. Therefore, we can conclude that the annotators had conducted high-quality labeling work. The disagreement is mainly caused by the non-trivial understanding of the tweets themselves. Such challenges in real-world tweets sentiment annotation also make our competition extremely meaningful. We expect the advanced models developed by the participants can be the effective solutions to this problem.

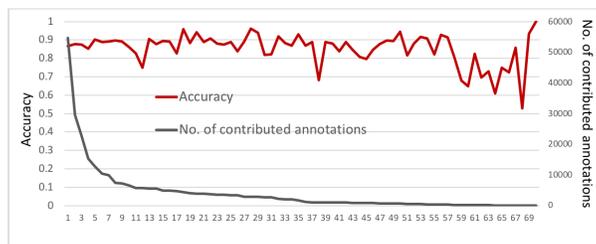


Figure 3: The accuracy of 69 annotators, ordered by their number of contributed annotations.

Table 3: Class distribution in data splits. The *All* set represents the complete collected dataset. *TRAINING*, *TEST1* and *TEST2* are the three splits used in the competition and also in the benchmark models evaluation. The three splits maintained the same class distribution as the original complete dataset.

	TRAINING		TEST1		TEST2		All	
	No. tweets	(%)						
Positive	8821	0.16	3150	0.16	3244	0.16	15215	0.16
Negative	8820	0.16	3252	0.16	3195	0.16	15267	0.16
Neutral	37359	0.68	13598	0.68	13561	0.68	64518	0.68
Total	55000	1.00	20000	1.00	20000	1.00	95000	1.00

2.4 Corpus Properties

ASAD consists of 95,000 annotated tweets, with three-class sentiments. Table 3 presents basic statistics about ASAD. ASAD (represented by ALL in the table) has a total of 15,215 *positive* tweets, 15,267 *negative* tweets and 64,518 *neutral* tweets. In the Kaggle competition, the entire ASAD data is split into *TRAINING*, *TEST1* and *TEST2* for the usage at different stages. The data distribution in all subset splits are the same. We also use these data split for benchmark evaluation, which is described in Section 3.

3 Benchmark Evaluation

This section provides details on evaluation details and baseline results.

3.1 Competition and Evaluation Details

In the competition, the ASAD dataset is divided into three folds: training (TRAINING), development-time testing (TEST1) and main testing (TEST2). The TRAINING set contains 55K labeled tweets, while TEST1 and TEST2 are composed of 20K tweets each, as shown in Table 3. In the competition, both TRAINING and TEST1 datasets are released at the beginning of the competition. At this stage, competing teams can assess their relative rank using the evaluation metric described next on the TEST1 set. At the end of the competition, teams will be ranked based on their performance on TEST2. Teams will be ranked on TEST2, where ranking will be done offline by evaluating the submitted codes. The prizes for the first three winners are:

- First place winner, 10,000 USD
- Second place winner, 5,000 USD
- Third place winner, 2,000 USD

At the end of the competition, a report will be published by inviting the top ranking teams to describe their winning solutions. And an awarding conference will be organized to invite the team introducing their solutions. For more details about the competition, timelines and rules please visit the competition website: <https://wti.kaust.edu.sa/arabic-sentiment-analysis-challenge>.

Given the highly imbalanced dataset and the multi-class nature of the problem and following [24, 28], we use the following metrics in the competition:

- **Primary metric**, the macro-averaged recall (recall averaged across the three classes)

$$AvgRec = \frac{1}{3}(R^P + R^N + R^U) \quad (4)$$

where R^P , R^N and R^U are the recall for the positive, negative and neutral class, respectively.

- **Secondary metric**, the average F_1 for positive and negative classes only

$$F_1^{PN} = \frac{1}{2}(F_1^P + F_1^N) \quad (5)$$

where F_1^P and F_1^N refer to the $F1$ measure for the positive and negative classes, respectively. For more details on the evaluation metric, please refer to [29].

The two proposed evaluation metrics take into consideration the multi-class sentiment classification problem, as well as the imbalanced nature of the dataset. On the released competition, participants will be evaluated and ranked by their performance on the primary metric only. The secondary metric is introduced to provide additional analysis of the results, and will not be used while ranking participating teams.

3.2 Our Baselines and Results

Preprocessing: In our baseline experiments we applied some preprocessing to the data by removing the stop-words, special symbols like '@' and '#', and URLs. We also removed non-Arabic characters in order to make a clear analysis of Arabic language. In addition, the word tokenization was applied to the text data. For this part we used the NLTK (<https://www.nltk.org>) and pyarabic tool (<https://github.com/linuxscout/pyarabic>).

Traditional Baseline Methods: We applied bag-of-words (BOW) and Tf-Idf to extract the features of the text data, then for each of them we classify by using logistic regression.

- **BOW+ Logistic regression.** The BOW is the simplest form of text representation, which describes the occurrence of words within a document. We used BOW to extract the features of the texts, then we fed them into the logistic regression classifier with 'lbfgs' solver.
- **TF-IDF + Logistic regression.** TF-IDF measures how important a particular word is with respect to a document and the entire corpus. Similar to the previous model, the text features were extracted by Tf-Idf, then we applied the logistic regression classifier.

For BOW, TF-Idf, and logistic regression, we used the popular Scikit-learn⁸ python library. Scikit-learn is an open source machine learning library that supports supervised and unsupervised learning. It also provides various tools for model fitting, data preprocessing, model selection and evaluation, and many other utilities.

Deep Learning Baseline Models: We fine-tuned the pre-trained model BERT [30] and AraBERT [31] for learning the embedding vectors for tweets and then for the classification.

- **BERT** [30] is one of the most popular language representation models developed by Google. BERT can be fine-tuned with one additional output layer to create a model for a specific NLP task. Several multilingual BERT models have been released to support many languages.
- **AraBERT** [31] is an Arabic language model by training BERT on a large-scale Arabic corpus. It overcomes the performance of a multilingual version of BERT in many NLP tasks applied to Arabic text.

To fine-tune BERT and AraBERT, we used ktrain⁹, which is a lightweight wrapper for many deep learning libraries to help build, train, and deploy machine learning models in a more accessible and easier manner. With ktrain, the text must be pre-processed in a specific way to obtain a fixed-size sequence of word IDs for each document. Then, they are used as input for the classifier. Each model in ktrain has its corresponding pre-processing procedure and classifier.

Table 4 presents the results of the baseline models when trained by TRAINING and evaluated on TEST1 and TEST2. These initial results provide a baseline for the competing teams, where a significant improvement is expected. Additionally, it is clear that the evaluation results on TEST1 is statistically similar to those obtained when testing the same model on TEST2. This indicates that the relative ranking on TEST1 is a good indicator for the participating teams of their final ranking on TEST2. Comparing the results of different models, we see AraBert has the best performance in terms of $Avg - Rec$ and F_1^{PN} , as well as the general metrics Macro-F1 and Micro-F1.

⁸https://scikit-learn.org/stable/getting_started.html

⁹<https://github.com/amaiya/ktrain>

Table 4: Evaluation results of baseline models on TEST1 and TEST 2

	BOW-LR		TF-IDF-LR		BERT		AraBERT	
	TEST1	TEST2	TEST1	TEST2	TEST1	TEST2	TEST1	TEST2
F1-Positive	0.51	0.51	0.52	0.53	0.62	0.61	0.65	0.65
F1-Negative	0.49	0.48	0.51	0.50	0.39	0.38	0.53	0.52
F1-Neutral	0.83	0.83	0.83	0.83	0.85	0.85	0.86	0.87
Macro-F1	0.61	0.60	0.62	0.62	0.62	0.62	0.68	0.68
Micro-F1	0.72	0.72	0.73	0.73	0.74	0.74	0.78	0.78
Positive-Recall	0.43	0.43	0.47	0.47	0.60	0.59	0.64	0.63
Negative-Recall	0.43	0.42	0.47	0.45	0.28	0.27	0.45	0.43
Neutral-Recall	0.88	0.88	0.87	0.87	0.92	0.92	0.90	0.91
$Avg - Rec$	0.58	0.58	0.60	0.60	0.60	0.59	0.66	0.66
F_1^{PN}	0.50	0.50	0.52	0.52	0.50	0.50	0.59	0.59

4 Discussion

In this section, we highlight some of the challenges related to Arabic sentiment classification from a Twitter-based dataset. We identify three main challenges, listed below:

- **Sentiment Imbalance:** ASAD adopts a three-point sentiment scale, where a tweet is labeled as either *positive*, *negative* or *neutral*. As it was detailed in Section 2, the distribution of the labels is imbalanced, where the *neutral* class has more number of tweets than the other two classes. Imbalanced datasets create a challenge to most learning algorithms. Thus, the prediction performance of the model should be evaluated using a metric that handles unbalanced data.
- **Spam:** The ASAD dataset was preliminary cleaned as discussed in Section 2. However, spam tweets still exist and may affect the prediction performance of a sentiment classifier.
- **Multiple Dialects:** The existence of multiple dialects in the dataset creates a challenge to the sentiment classification problem. In a simpler dataset with one dialect only, e.g., ASTD [25], a supervised learning algorithm will be able to identify the main keywords associated with each sentiment in this dialect. When multiple dialects exists, the range of keywords associated with each sentiment increases because such words differ from one dialect to another. This creates an additional challenge that one-dialect corpus does not have. A pre-trained word embedding model with one-dialect corpus may not work well on the multi-dialect dataset.

The complete dataset with both training and testing sets will be publicly released for the research community after the competition ends. The dataset can be used for a range of other NLP tasks, besides sentiment classification. This includes dialect identification, spam detection, and also assist the sentiment analysis of Covid-19 tweets [32]. We will update our work in the near future.

References

- [1] Imane Guellil, Houda Saâdane, Faical Azouaou, Billel Gueni, and Damien Nouvel. Arabic natural language processing: An overview. *Journal of King Saud University-Computer and Information Sciences*, 2019.
- [2] Mahmoud Al-Ayyoub, Abed Allah Khamaiseh, Yaser Jararweh, and Mohammed N Al-Kabi. A comprehensive survey of arabic sentiment analysis. *Information processing & management*, 56(2):320–342, 2019.
- [3] Noura Farra, Elie Challita, Rawad Abou Assi, and Hazem Hajj. Sentence-level and document-level sentiment mining for arabic texts. In *2010 IEEE international conference on data mining workshops*, pages 1114–1119. IEEE, 2010.
- [4] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [5] Ahmad Al Sallab, Hazem Hajj, Gilbert Badaro, Ramy Baly, Wassim El-Hajj, and Khaled Shaban. Deep learning models for sentiment analysis in arabic. In *Proceedings of the second workshop on Arabic natural language processing*, pages 9–17, 2015.
- [6] Sadam Al-Azani and El-Sayed M El-Alfy. Hybrid deep learning for sentiment polarity determination of arabic microblogs. In *International Conference on Neural Information Processing*, pages 491–500. Springer, 2017.

- [7] Maha Heikal, Marwan Torki, and Nagwa El-Makky. Sentiment analysis of arabic tweets using deep learning. *Procedia Computer Science*, 142:114–122, 2018.
- [8] Mohammad Al-Smadi, Omar Qawasmeh, Mahmoud Al-Ayyoub, Yaser Jararweh, and Brij Gupta. Deep recurrent neural network vs. support vector machine for aspect-based sentiment analysis of arabic hotels’ reviews. *Journal of computational science*, 27:386–393, 2018.
- [9] Mahmoud Al-Ayyoub, Aya Nuseir, Kholoud Alsmearat, Yaser Jararweh, and Brij Gupta. Deep learning for arabic nlp: A survey. *Journal of computational science*, 26:522–531, 2018.
- [10] Ammar Mohammed and Rania Kora. Deep learning approaches for arabic sentiment analysis. *Social Network Analysis and Mining*, 9(1):52, 2019.
- [11] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [12] Abu Bakr Soliman, Kareem Eissa, and Samhaa R El-Beltagy. Aravec: A set of arabic word embedding models for use in arabic nlp. *Procedia Computer Science*, 117:256–265, 2017.
- [13] Obeida ElJundi, Wissam Antoun, Nour El Droubi, Hazem Hajj, Wassim El-Hajj, and Khaled Shaban. hulmona: The universal language model in arabic. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 68–77, 2019.
- [14] Mahmoud Nabil, Mohamed Aly, and Amir Atiya. Labr: A large scale arabic sentiment analysis benchmark. *arXiv preprint arXiv:1411.6718*, 2014.
- [15] Muhammad Abdul-Mageed and Mona T Diab. Sana: A large scale multi-genre, multi-dialect lexicon for arabic subjectivity and sentiment analysis. In *LREC*, pages 1162–1169, 2014.
- [16] Mohamed Aly and Amir Atiya. Labr: A large scale arabic book reviews dataset. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 494–498, 2013.
- [17] Hady ElSahar and Samhaa R El-Beltagy. Building large arabic multi-domain resources for sentiment analysis. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 23–34. Springer, 2015.
- [18] Muhammad Abdul-Mageed and Mona T Diab. Awatif: A multi-genre corpus for modern standard arabic subjectivity and sentiment analysis. In *LREC*, volume 515, pages 3907–3914, 2012.
- [19] Joseph L Fleiss. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378, 1971.
- [20] Housseem Abdellaoui and Mounir Zrigui. Using tweets and emojis to build tead: an arabic dataset for sentiment analysis. *Computación y Sistemas*, 22(3):777–786, 2018.
- [21] Kathrein Abu Kwaik, Stergios Chatzikyriakidis, Simon Dobnik, Motaz Saad, and Richard Johansson. An arabic tweets sentiment analysis dataset (atsad) using distant supervision and self training. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 1–8, 2020.
- [22] Nora Al-Twairish, Hend Al-Khalifa, AbdulMalik Al-Salman, and Yousef Al-Ohali. Arasenti-tweet: A corpus for arabic sentiment analysis of saudi tweets. *Procedia Computer Science*, 117:63–72, 2017.
- [23] Eshrag Refaee and Verena Rieser. An arabic twitter corpus for subjectivity and sentiment analysis. In *LREC*, pages 2268–2273, 2014.
- [24] Sara Rosenthal, Noura Farra, and Preslav Nakov. Semeval-2017 task 4: Sentiment analysis in twitter. In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, pages 502–518, 2017.
- [25] Mahmoud Nabil, Mohamed Aly, and Amir Atiya. Astd: Arabic sentiment tweets dataset. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 2515–2519, 2015.
- [26] Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. Semeval-2018 Task 1: Affect in tweets. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA, 2018.
- [27] Victoria Bobicev and Marina Sokolova. Inter-annotator agreement in sentiment analysis: Machine learning perspective. In *RANLP*, pages 97–102, 2017.
- [28] Fabrizio Sebastiani. An axiomatically derived measure for the evaluation of classification algorithms. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*, pages 11–20, 2015.
- [29] Preslav Nakov, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. Semeval-2016 task 4: Sentiment analysis in twitter. *arXiv preprint arXiv:1912.01973*, 2019.

- [30] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [31] Wissam Antoun, Fady Baly, and Hazem Hajj. Arabert: Transformer-based model for arabic language understanding, 2020.
- [32] Qiang Yang, Hind Alamro, Somayah Albaradei, Adil Salhi, Xiaoting Lv, Changsheng Ma, Manal Alshehri, Inji Jaber, Faroug Tifratene, Wei Wang, Takashi Gojobori, Carlos M. Duarte, Xin Gao, and Xiangliang Zhang. SenWave: Monitoring the global sentiments under the covid-19 pandemic. *arXiv preprint arXiv:2006.10842*, 2020.