

Improving Natural Language Inference in Arabic using Transformer Models and Linguistically Informed Pre-Training

1st Mohammad Majd Saad Al Deen[§]

Fraunhofer IAIS and Hochschule Bonn-Rhein-Sieg
Sankt Augustin, Germany

2nd Maren Pielka[§]

Fraunhofer IAIS
Sankt Augustin, Germany
maren.pielka@iais.fraunhofer.de

3rd Jörn Hees

Hochschule Bonn-Rhein-Sieg
Sankt Augustin, Germany

4th Bouthaina Soulef Abdou

Fraunhofer IAIS and University of Bonn
Sankt Augustin, Germany

5th Rafet Sifa

Fraunhofer IAIS and University of Bonn
Sankt Augustin, Germany

Abstract—This paper addresses the classification of Arabic text data in the field of Natural Language Processing (NLP), with a particular focus on Natural Language Inference (NLI) and Contradiction Detection (CD). Arabic is considered a resource-poor language, meaning that there are few data sets available, which leads to limited availability of NLP methods. To overcome this limitation, we create a dedicated data set from publicly available resources. Subsequently, transformer-based machine learning models are being trained and evaluated. We find that a language-specific model (AraBERT) performs competitively with state-of-the-art multilingual approaches, when we apply linguistically informed pre-training methods such as Named Entity Recognition (NER). To our knowledge, this is the first large-scale evaluation for this task in Arabic, as well as the first application of multi-task pre-training in this context.

I. INTRODUCTION

Natural Language Processing (NLP) in Arabic, also known as Arabic NLP, is a subfield of Artificial Intelligence (AI) that focuses on processing and analysing textual data in the Arabic language. It encompasses various technologies and methods for automating tasks such as text classification, sentiment analysis, and machine translation. The goal is to teach computers to understand and process Arabic language, enabling a range of applications including chatbots, text mining tools, and translation services. This is a challenging field due to the limited availability of training data and pre-trained models for the Arabic language.

Natural Language Inference (NLI), also referred to as Recognizing Textual Entailment (RTE), is a subfield of Natural Language Processing (NLP) that aims to determine the semantic relationship between two pieces of text, known as the “premise” and the “hypothesis”. The task, as described by MacCartney et al. [1], involves identifying possible connections such as “entailment” (if the premise is true, the hypothesis must also be true), “contradiction” (if the premise is true, the hypothesis cannot be true, and vice versa), or “neutral” (there is no logical relationship between the two sentences;

both can be true or false simultaneously). Accomplishing this task requires a machine learning algorithm to comprehend the semantics of a text, which poses a particularly challenging problem.

In this work, we apply existing deep learning approaches, namely AraBERT and XLM-RoBERTa, to the task of Natural Language Inference and Contradiction Detection in Arabic. To our best knowledge, this has not been attempted before in such a comprehensive manner, as previous studies were limited to a smaller number of data sources and models. We also employ an informed language modeling approach, by further pre-training the transformer model on an NER task, before fine-tuning it on the downstream tasks. This is a novel direction of research with respect to Arabic text. In addition, we introduce a new data collection for NLI/CD in Arabic language, which is publicly available on Github¹.

RELATED WORK

In the field of Natural Language Inference (NLI) and considering its impact on Question Answering (QA) tasks, Mishra et al. [3] conducted a research study. They utilized the RACE dataset [4], which is a large-scale reading comprehension dataset consisting of questions and answers from English exams for Chinese students. The authors converted a subset of the RACE dataset (containing 48,890 training examples, 2,496 validation examples, and 2,571 test examples) into an NLI format and compared the performance of a state-of-the-art model, RoBERTa [23], in both formats. To convert a reading comprehension question into a Natural Language Inference (NLI) format, the question was used as the premise, and each answer option was paraphrased as individual hypotheses. The same model architecture, comprising a RoBERTa encoder and a two-layer feed-forward network as the classification head, was employed for both QA and NLI. The results showed that the NLI model outperformed

[§]Both authors contributed equally.

¹https://github.com/fraunhofer-iais/arabic_nlp/

the QA model on the subset of the RACE dataset. This can be attributed to the more natural form of the hypotheses in the NLI model compared to the combination of question and answer option in the QA model.

With respect to NLI in Arabic, Jallad et al. [5] conducted a similar study and created their own dataset called arNLI, which consists of over 6,000 data points. The data was obtained by using machine translation from two English sources, namely SICK² and PHEME³. The authors developed a system with three main components: text preprocessing (cleaning, tokenization, and stemming), feature extraction (contradiction feature vector and language model vectors), and a machine learning classification model. The morphological units were processed using the Snowball Stemmer (Porter2) algorithm. Various types of features were employed for feature extraction, including features for named entities, similarity features, specific stopword features, number, date, and time features, which were processed using different language models such as TF-IDF [6] and Word2Vec embeddings [7]. To determine the relationship type between two sentences (contradiction, entailment, or neutral), various traditional machine learning classifiers were used, including Support Vector Machine (SVM), Stochastic Gradient Descent (SGD), Decision Tree (DT), ADA Boost Classifier, K-Nearest Neighbors (KNN), and Random Forest. The proposed solution by the authors was trained and evaluated on their own dataset (arNLI), and they reported that Random Forest achieved the best results on the arNLI dataset with an accuracy of 75%.

Contradiction detection has received relatively limited attention in the literature compared to other tasks. De Marneffe et al. [2] define a contradiction as a conflict between two statements that mutually negate each other. In a stricter logical sense, there is no possible world in which both statements can be simultaneously true. A looser definition, aligning better with human intuition, suggests that a contradiction exists when it is highly improbable for two statements to be simultaneously true. To identify these contradictions, the authors employ an approach based on syntactic analysis and semantic understanding of the text. They utilize syntactic parsing tools to establish the logical structure of the text, and semantic features such as antonymy, polarity, and numerical deviation to comprehend the meaning of words and sentences in the text. The authors note that detecting contradictions may be more challenging than recognizing entailment and requires a deep semantic understanding, possibly augmented by world knowledge.

Pucknat et al. [8] conducted a study comparing the performance of four state-of-the-art models in NLI, particularly for contradiction detection, on German text data. These models were evaluated based on their performance on a machine-translated version of the well-known Stanford Natural Language Inference dataset (SNLI) [9] and the German test set of the Cross-Lingual NLI Corpus (XNLI) [10]. One key focus was to determine if the models were robust with respect

to data selection and could potentially be applied in real-world scenarios. The XLM-RoBERTa model significantly outperformed the other models, likely due to its extensive pre-training and multi-head attention. However, the models did not generalize well to the XNLI data, indicating that the training corpus was limited in terms of topics and types of contradictions. The authors report an accuracy of 86.5% when testing XLM-RoBERTa on the XNLI dataset.

Another methodology by Pielka et al. [16] focused on pre-training methods to integrate syntactic and semantic information into state-of-the-art model architectures. The authors presented a linguistically enhanced approach for pre-training transformer models. They incorporated additional knowledge about part-of-speech tags, syntactic analysis, and semantic relationships between words into the model. Their goal was to become independent of massive pre-training data resources by integrating as much external knowledge as possible into the model. Their approach was evaluated on the SNLI dataset, and they demonstrated that the smaller BERT model can be competitive with XLM-RoBERTa when enhanced with additional knowledge during pre-training. Their approach did not require additional data for pre-training, as it was trained on additional tasks using the same dataset that would later be used for fine-tuning.

DATA

For this study, a self-constructed corpus was used, comprising data from three different sources.

- XNLI (Cross-Lingual NLI Corpus) [10]: The Arabic-translated section of the XNLI dataset was included as a source for the corpus. XNLI is a well-known benchmark dataset for cross-lingual natural language inference tasks, containing 7500 text pairs in 15 languages.
- SNLI (Stanford Natural Language Inference Corpus) [9]: The Arabic-translated section of the SNLI dataset by [11], comprising 1332 manually translated sentence pairs, was also incorporated into the corpus. SNLI is a widely used dataset in English language for NLI, consisting of sentence pairs labeled with three relationship types: entailment, contradiction, and neutral.
- arNLI (Arabic Natural Language Inference) [5]: The arNLI dataset, specifically created for the NLI task in Arabic language, was an additional source of data. This dataset consists of 6366 data points and was obtained through machine translation from two English sources.

Some examples from the data set are displayed in Table I.

As the data was compiled from different sources, necessary standard normalizations with respect to encoding, column names, label mappings etc. were performed.

In the context of using transformer-based models for NLI/CD tasks, additional preprocessing steps such as stemming or stopword removal are not necessary.

After performing the general preprocessing for the data from the three sources, the next step is to merge the data to create a unified dataset. To ensure a fair distribution of training, testing, and validation data, the merged data is being randomly shuffled before splitting. The split is then done with

²<https://alt.qcri.org/semeval2014/task1/>

³<https://www.kaggle.com/datasets/usharengaraju/pHEME-dataset>

premise	hypothesis	label	premise_en	hypothesis_en
لقد دخل أول فريق للتدخل السريع التابع لشرطة نيويورك ردهة الشارع الغربي للبرج الشمالي وأهم مستعدون لبدء التسلق حوالي ٩:١٥ صباحاً	كان البرج لا يزال قائماً في الساعة ٩:١٥ صباحاً	0	The first NYPD ESU team entered the West Street-level lobby of the North Tower and prepared to begin climbing at about 9:15 A.M	The tower was still standing at 9:15 AM.
يتطلب الأمر مشاركة بين الدعم الخاص والتمويل الجامعي لدرستنا القانونية لمواصلة النمو في المكانة والثأني	مدرسة القانون لدينا مدعومة جزئياً بواسطة مؤسسة ميلندا وبيل جيت	1	It takes a partnership of private support and University funding for our law school to continue to grow in stature and influence	Our law school is supported in part by the Melinda and Bill Gates Foundation
يجب على الأمريكيين أيضاً أن يفكروا في كيفية القيام بذلك وتنظيم حكومتهم بطريقة مختلفة	يمكن تنظيم الحكومة فقط بطريقة واحدة وأي محاولة لتغييرها ستكون غبية	2	Americans should also consider how to do it-organizing their government in a different way.	The government can only be organized in one way and any attempt to change it would be foolish
البوابة التي تمثل جزءاً من جدار المدينة، لم يكن المقصود منها من البروسيين الأكثر براغماتية أكثر من مجرد قوس النصر كمعبر لفرض الرسوم	كانت البوابة مجرد قوس نصر.	2	Forming part of the city wall, the gate was intended by the more pragmatic Prussians not so much as a triumphal arch as an imposing tollgate for collecting duties.	The gate was just a triumphal arch

TABLE I

FOUR EXAMPLES SHOWING THE ARABIC TEXT DATA, ENGLISH TRANSLATION AND LABEL (0: ENTAILMENT, 1: NEUTRAL, 2: CONTRADICTION).

a distribution of 60% for training data, 20% for testing data, and 20% for validation data.

The final data set consists of a total of 14,758 pairs of premises and hypotheses, each accompanied by an English translation and a label. The labels are encoded as 0 (entailment), 1 (neutral), and 2 (contradiction). Retaining the English translation of the data can assist in better comparing and understanding different texts and contexts, facilitating the comparison of various models.

For the NER pre-training, the ANERcorp corpus from CamelLabSplits [12] is being used, which contains 3973 text samples in Arabic language with annotations for the NER task. The entity types “person”, “location”, “organization” and “miscellaneous” are being used and annotated according to the IOB-scheme.

METHODOLOGY

In the scope of this study, analyses are conducted for Natural Language Inference (NLI) and Contradiction Detection (CD). The data processing pipeline remains the same for both tasks since the input data consists of the “premise” and “hypothesis” columns in both cases. However, differences arise in terms of the labels used. The labels of the datasets for the CD task are modified to treat the problem as a binary

classification. Originally, “0” was used for entailment, “1” for neutral, and “2” for contradiction. However, since only a binary outcome is required in case of CD, “0” and “1” are mapped to “0”, meaning there is no contradiction, and “2” to “1”, meaning there is a contradiction.

In this study, two state-of-the-art models, namely AraBERT [13] and XLM-RoBERTa [14], are being investigated. AraBERT is based on the BERT [15] paradigm and pre-trained on 24GB of news corpora in Arabic language with the Masked Language Modeling (MLM) and Next Sentence Prediction (NSP) objectives. A special sub-word segmentation algorithm is used to account for the semantic granularity of the Arabic language. For this study, the base version of AraBERT with 136M parameters is being used. XLM-RoBERTa is a multi-lingual language model, which has been pre-trained for MLM on 2.5TB of CommonCrawl data in 100 languages. It is shown to achieve state-of-the-art results on many NLP tasks across multiple languages. We use the base version with 279M parameters. To this end, we conduct the following modeling and evaluation steps:

- Loading the pre-trained model: We utilize the checkpoints “xlm-roberta-base”⁴ and “aubmindlab/bert-base-

⁴<https://huggingface.co/xlm-roberta-base>

arabertv02”⁵ for XLMRoBERTa and AraBERT, respectively.

- Additional pre-training: We employ Named Entity Recognition (NER) as an additional pre-training step before fine-tuning the model on the downstream task. Following the findings from [16], the idea is to provide the transformer with semantic knowledge that would help in identifying contradictions and entailments. In the pre-training phase, a word-level classification head is being attached to the encoder and later discarded.
- Model fine-tuning: The pre-trained model is then fine-tuned on the tasks (NLI and CD) using the created datasets.
- Hyperparameter optimization is performed during the finetuning process to enhance the model’s performance. Techniques such as learning rate scheduling, dropout, and batch size adjustment are employed.
- Model evaluation: After model fine-tuning, the performance of the model is evaluated using metrics such as accuracy or F1 score.

The final preparation of input data involves merging the “premise” and “hypothesis” columns in the training, testing, and validation datasets using special tokens ([CLS] and [SEP]). The separator token “[SEP]” is recognized by the BERT tokenizer as a special token to separate different parts of the text. This allows the BERT model to process the meaning of each text part separately and then combine them in a single step. The classification token “[CLS]” in the BERT architecture signals the model to perform classification on the input text. It is placed at the beginning of the text and serves as an indication for the model to make predictions about the text’s class membership.

Overall, this study demonstrates the effective utilization of transformer-based models, showcasing their adaptability and performance in NLI and CD tasks in the Arabic language.

II. EXPERIMENTS AND RESULTS

We mainly compare AraBERT and XLM-RoBERTa with respect to their performance on the NLI and CD task. An extensive hyperparameter search is being conducted, using the Optuna [17] tool with 150 optimization runs. We train the models for a maximum of 5 epochs, with the option of early stopping if the validation performance does not improve any further. The AdamW optimizer [18] is being applied. The other hyperparameters, including learning rate, weight decay and batch size, are chosen according to the best result from the respective Optuna run. The experimental results on the two tasks are displayed in tables II and III.

All the models achieve overall good results on both tasks. It is especially noteworthy that AraBERT performs competitively with XLM-RoBERTa, even though it was pre-trained on a considerably smaller amount of data. With respect to the CD task, the best AraBERT model with multitask finetuning even outperforms XLM-RoBERTa by two percentage points. This emphasizes the fact that language-specific finetuning can

Model	Multitask Finetuning	Accuracy	F1-Score
AraBERT	x	75.3	75.4
XLM-R	x	78.7	78.8
AraBERT	✓	76.8	76.8
XLM-R	✓	78.9	79.0

TABLE II
RESULTS FOR THE NLI TASK WITH ARABERT & XLM-ROBERTA, IN %. ACCURACY AND MACRO AVERAGE F1-SCORE ARE BEING REPORTED.

Model	Multitask Finetuning	Accuracy	F1-Score
AraBERT	x	87.4	82.3
XLM-R	x	86.9	81.0
AraBERT	✓	88.1	84.5
XLM-R	✓	86.8	81.1

TABLE III
RESULTS FOR THE CD TASK WITH ARABERT & XLM-ROBERTA, IN %. ACCURACY AND MACRO AVERAGE F1-SCORE ARE BEING REPORTED.

be more effective than extensive multi-lingual pre-training for some downstream tasks. We also find that adding the NER objective as an additional pre-training step improves the performance. Interestingly, this effect is stronger for the smaller AraBERT model, suggesting that it can help bridge the performance gap that is caused by XLMRoBERTa’s larger model size and the amount of training data it has access to.

III. CONCLUSION AND OUTLOOK

We presented the first comprehensive study on Natural Language Inference and Contradiction Detection in Arabic language, in which we applied state-of-the-art transformer methods combined with an informed pre-training approach. The methods achieve promising results on our custom data set, emphasizing the fact that smaller, language-specific models like AraBERT can perform competitively with larger multi-lingual models, if they are being enhanced with additional linguistic knowledge. Further, we collected a large data set for NLI in Arabic language.

Future work includes adding more pre-training methods such as Part of Speech tagging, Word Sense Disambiguation or Semantic Role Labeling. We expect the performance of the AraBERT model to improve even further by adding more linguistic knowledge. Another direction of research is to exploit the potential of large language models such as GPT-4 [19], by casting the classification problem as a text generation task. It would be interesting to see the performance of those resourceful models when confronted with a low-resource language such as Arabic.

ACKNOWLEDGEMENTS

This research has been partially funded by the Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as part of the Lamarr-Institute for Machine Learning and Artificial Intelligence.

⁵<https://huggingface.co/aumindlab/bert-base-arabertv02>

REFERENCES

- [1] MacCartney, Bill and Manning, Christopher: Natural Logic for Textual Inference, 2007. doi : 10.3115/1654536.1654575.
- [2] M.-C. de Marneffe, A. N. Rafferty und C. D. Manning. Finding Contradictions in Text. In: Proceedings of ACL-08: HLT. Columbus, Ohio: Association for Computational Linguistics, 2008. url: <https://aclanthology.org/P08-1118>.
- [3] A. Mishra, D. Patel, A. Vijayakumar, X. Li, P. Kapanipathi und K. Talama-dupula. Reading Comprehension as Natural Language Inference: A Semantic Analysis. In: Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics. Barcelona, Spain (Online): Association for Computational Linguistics, 2020. url: <https://aclanthology.org/2020.starsem-1.2>.
- [4] G. Lai, Q. Xie, H. Liu, Y. Yang und E. Hovy. RACE: Large-scale Reading Comprehension Dataset From Examinations. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. Copenhagen, Denmark: Association for Computational Linguistics, 2017. url: <https://aclanthology.org/D17-1082>.
- [5] K. A. Jallad und N. Ghneim. ArNLI: Arabic Natural Language Inference for Entailment and Contradiction Detection. In: arXiv preprint arXiv:2209.13953, 2022.
- [6] K. Sparck Jones: A statistical interpretation of term specificity and its application in retrieval. en. In: Journal of documentation 28.1., 1972.
- [7] T. Mikolov, K. Chen, G. Corrado und J. Dean: Efficient estimation of word representations in vector space. In: arXiv preprint arXiv:1301.3781, 2013.
- [8] L. Pucknat, M. Pielka und R. Sifa. Detecting Contradictions in German Text: A Comparative Study. In: 2021 IEEE Symposium Series on Computational Intelligence (SSCI), 2021. doi: 10.1109/SSCI50451.2021.9659881.
- [9] S. R. Bowman, G. Angeli, C. Potts und C. D. Manning. A large annotated corpus for learning natural language inference. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. Lisbon, Portugal: Association for Computational Linguistics, 2015. url: <https://aclanthology.org/D15-1075>.
- [10] A. Conneau, R. Rinott, G. Lample, A. Williams, S. Bowman, H. Schwenk und V. Stoyanov. XNLI: Evaluating Cross-lingual Sentence Representations. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. Brussels, Belgium: Association for Computational Linguistics, 2018. url: <https://aclanthology.org/D18-1269>.
- [11] Ž. Agić und N. Schluter: Baselines and Test Data for Cross-Lingual Inference. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018), 2018.
- [12] O. Obeid, N. Zalmout, S. Khalifa, D. Taji, M. Oudah, B. Alhafni, G. Inoue, F. Eryani, A. Erdmann und N. Habash. CAMEL Tools: An Open Source Python Toolkit for Arabic Natural Language Processing. In: Proceedings of the Twelfth Language Resources and Evaluation Conference. Marseille, France: European Language Resources Association, 2020.
- [13] W. Antoun, F. Baly und H. Hajj: AraBERT: Transformer-based Model for Arabic Language Understanding, 2020. url: <https://arxiv.org/abs/2003.00>.
- [14] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer und V. Stoyanov: Unsupervised Cross-lingual Representation Learning at Scale, 2019. url: <https://arxiv.org/abs/1911.02>.
- [15] J. Devlin, M.-W. Chang, K. Lee und K. Toutanova: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: arXiv preprint arXiv:1810.04805, 2018.
- [16] M. Pielka, S. Schmidt, L. Pucknat und R. Sifa. Towards Linguistically Informed Multi-Objective Pre-Training for Natural Language Inference, 2022. arXiv: 22.12.07428 [cs.CL].
- [17] T. Akiba, S. Sano, T. Yanase, T. Ohta und M. Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework. arXiv: 1907.10902 [cs.CL], 2019.
- [18] I. Loshchilov und F. Hutter: Decoupled Weight Regularization. In: Proceedings of ICLR 2019, 2019.
- [19] OpenAI: GPT-4 Technical Report. arXiv: 2303.08774, 2023.