

Context-Gloss Augmentation for Improving Arabic Target Sense Verification

Sanad Malaysha, Mustafa Jarrar, Mohammed Khalilia

Birzeit University, Palestine

{smalaysha, mjarrar, mkhalilia}@birzeit.edu

Abstract

Arabic language lacks semantic datasets and sense inventories. The most common semantically-labeled dataset for Arabic is the *ArabGlossBERT*, a relatively small dataset that consists of 167K context-gloss pairs (about 60K positive and 107K negative pairs), collected from Arabic dictionaries. This paper presents an enrichment to the ArabGlossBERT dataset, by augmenting it using (Arabic-English-Arabic) machine back-translation. Augmentation increased the dataset size to 352K pairs (149K positive and 203K negative pairs). We measure the impact of augmentation using different data configurations to fine-tune BERT on target sense verification (TSV) task. Overall, the accuracy ranges between 78% to 84% for different data configurations. Although our approach performed at par with the baseline, we did observe some improvements for some POS tags in some experiments. Furthermore, our fine-tuned models are trained on a larger dataset covering larger vocabulary and contexts. We provide an in-depth analysis of the accuracy for each part-of-speech (POS).

1 Introduction

There are three tasks in the literature that are related to semantic understanding of natural language: (i) Word Sense Disambiguation (WSD), (ii) Target Sense Verification (TSV), and (iii) Word-in-Context (WiC). WSD is the most common task, which aims to disambiguate word's semantics. Given a context (i.e., sentence), a target word in the context, and a set of candidate senses (i.e., glosses, meaning definitions (Jarrar, 2006)) for the target word, the goal of the WSD task is to determine *which of these senses* is the intended meaning for the target word (Al-Hajj and Jarrar, 2022). For example, the word (جداول *ġdāwl*) has two senses in Arabic: *tables* (شكل يحتوي على مجموعة قضايا أو معلومات) and *creek* (مجرى صغير متفرع من نهر). Thus, in the context (تمشى بين الجداول والازهار), WSD aims to determine which of the two senses is the intended meaning of (الجداول *alġdāwl*). The TSV task is newly proposed in the literature (Breit et al., 2020). It aims to classify a sentence pair with positive or negative. Given a context, target and gloss, TSV aims to decide

whether this gloss is the intended meaning of the target. In other words, TSV does not determine which sense is the intended meaning, but rather, decides whether the context-gloss pair match (Positive) or not (Negative). For example, the sentence pair (تمشى بين الجداول والازهار - مجرى صغير متفرع من نهر) is Positive, while the pair (شكل يحتوي على مجموعة قضايا أو معلومات - تمشى بين الجداول والازهار) is Negative. WiC aims to determine whether a target word in two contexts is used in the same sense or not (Moreno et al., 2021). Although the three tasks are closely related, they are not the same, and the choice of which task to use depends on the application scenario (e.g., machine translation, information retrieval, semantic tagging, or others). Some researchers try to address these tasks using different approaches. For example, Hauer and Kondrak (2022) proposed to solve the WSD by re-formulating it as a TSV task, a WiC task and a combination of TSV and WiC tasks.

Such semantic understanding tasks have been challenging for many years, but recently gained attention due to the advances in contextualized word embedding models (Al-Hajj and Jarrar, 2022, 2021). Language models, specially BERT (Kenton and Toutanova, 2019), have made significant advancements in down-streaming NLP tasks. BERT is a transformer-based model pre-trained on huge corpora (Devlin et al., 2019). It can be fine-tuned on domain/task-specific data (e.g., POS tagging, WSD, TSV, and WiC) to update its contextualized embeddings. The TSV task has been addressed by fine-tuning BERT on context-gloss pairs as a sentence pair binary classification problem (Huang et al., 2019; Yap et al., 2020; Patel et al., 2021; Ranjbar and Zeinali, 2021; Lin and Giambi, 2021; El-Razzaz et al., 2021; Al-Hajj and Jarrar, 2022). However, the TSV task, similar to most NLP tasks, suffers from the knowledge-gain bottleneck, i.e., the lack of available quality datasets to train machine learning models.

Arabic is a low-resourced language (Darwish

et al., 2021; Jarrar et al., 2022) and the only available context-gloss pairs dataset is ArabGlossBERT (Al-Hajj and Jarrar, 2022). It consists of 167K context-gloss pairs, a relatively small dataset for fine-tuning BERT on a TSV task. The positive pairs (60K) were collected from multiple Arabic dictionaries (Jarrar and Amayreh, 2019; Jarrar, 2018; Jarrar et al., 2019; Jarrar, 2020) as well as from the Arabic Ontology (Jarrar, 2021, 2011). The pairs were cross-related to generate 106.8K negative pairs and used to fine-tune a BERT model, which achieved 84% accuracy.

This paper aims to enrich the ArabGlossBERT dataset by augmenting it using the back-translation technique, similar to the work done for English by Lin and Giambi (2021). With data augmentation, we generate new Arabic paraphrased contexts and glosses by translating the original data into English and back to Arabic, using Google Translate API. The new augmented dataset consists of 352K context-gloss pairs. To answer the question of whether the back-translation enrichment improves the TSV accuracy, we conduct 13 experiments that compare the accuracy obtained using the original dataset with the accuracy obtained using different combinations of the augmented datasets. We, also, provide an in-depth analysis of the TSV accuracy for each part-of-speech, which was not provided in (Al-Hajj and Jarrar, 2022). The main contributions of this work are:

- Augmented ArabGlossBERT using back-translation (352K pairs).
- Thirteen experiments with different dataset configurations - to measure whether the back-translation enrichment can improve TSV performance.
- In-depth analysis of the TSV accuracy for each part-of-speech.

The rest of the paper is organized as follows: Section 2 reviews the related literature, Section 3 describes the data augmentation, Section 4 presents the experiments, Section 5 presents the results and we conclude in Section 6 with limitations and future work.

2 Related Work

TSV has proven to be an effective solution for the WSD in many state-of-the-art efforts. Although

some researchers did not use the term TSV, this notion was also referred to as GlossBERT or *Context-Gloss Binary Classification* (Al-Hajj and Jarrar, 2022; El-Razzaz et al., 2021). A TSV training dataset is typically a set of context-gloss pairs, each labeled with Positive or Negative. A pre-trained language model can then be fine-tuned for sentence pair binary classification. This idea was first proposed for English as GlossBERT (Huang et al., 2019), where the training pairs were generated from a known SemCor dataset (Miller et al., 1993) with the target word, in context, marked up by a BERT-specific signal to emphasize it in the learning phase.

A similar effort in (Lin and Giambi, 2021) followed the GlossBERT technique. Their addition is the use of back-translation for improving the English WSD. They used back-translation from English to German and back to English in order to bridge the knowledge-gain gap and provide more context-gloss pairs. They also used a mark-up signal to surround the target word with double quotations. Only 2% improvement was achieved using back-translation. This paper aims to evaluate this idea for Arabic. Another idea was proposed in (Yap et al., 2020), in which a learning signal (special token [TGT]) was used, and BERT was fine-tuned on sequence-pair ranking, the model selects the most related gloss given a context sentence and a list of candidate glosses. Botha et al. (2020) used different mark-up signals in the form of open and close tags to emphasize the target word [E]target[/E] within the context sentence.

For Arabic, the TSV task was addressed in (Al-Hajj and Jarrar, 2022), which presents the ArabGlossBERT, a dataset of 167K context-gloss pairs labeled with Positive or Negative. First, 60K positive pairs were extracted from different Arabic lexicons, then 106K negative pairs were generated automatically by cross-relating the positive pairs. The target word was marked-up with different learning signals. Different Arabic pre-trained models were fine-tuned, and the best model using AraBERT-V2 (Antoun et al., 2020) achieved 84% accuracy. Similar work for Arabic was proposed in (El-Razzaz et al., 2021) using a smaller dataset (15K positive and 15K negative) in which they used AraBERT-V2 and reported 89% F1-score but this performance was criticized in (Al-Hajj and Jarrar, 2022).

	Gloss	-	Context
Original (Arabic)	فكرة أو مسألة تقدم للبحث	-	جلس المسؤولون يناقشون أطروحات المشروع
Translated (Arabic to English)	An idea or question is progressing to research	-	Officials sat discussing project proposals
Back-Translated (English to Arabic)	فكرة أو سؤال مقدم للبحث	-	جلس المسؤولون لمناقشة مقترحات المشاريع

Table 1: Example of context-gloss back-translation (Arabic-English-Arabic).

3 Data Augmentation

NLP tasks, including TSV, typically suffer from knowledge acquisition. The importance of knowledge acquisition is increasing especially because most NLP tasks are currently tackled using pre-trained neural models such as BERT, which generally requires large data to fine-tune. If the training data is not sufficient, the model will encounter the problem of unseen vocabulary and contexts, which harms model accuracy (Bevilacqua et al., 2021). The linguistic resources that can be utilized for semantic understanding tasks are limited in Arabic language. Our assumption, for the TSV task, is that the more context-gloss pairs can be used during the training phase, the more vocabulary and more contexts will be covered, thus the better TSV accuracy. This is why researchers started to try new techniques for data augmentation in order to enrich the available dataset with more knowledge (Lin and Giambi, 2021; Ranjbar and Zeinali, 2021).

For Arabic, and in order to enrich existing Arabic datasets, we propose to use the Arabic-English-Arabic back-translation, as illustrated in Table 1. It shows how the back-translation of glosses and contexts generates new paraphrased sentences with the same meaning. For back-translation we used Google Translate API, which was found to produce good quality and generally acceptable translations (De Vries et al., 2018). We did not remove diacritics since Arabic is diacritic-sensitive (Jarrar et al., 2018). Nevertheless, there are sentences that appeared with wrong or bad-quality translations, which we will discuss later. The translation was done in two phases. The glosses and contexts were translated into English, then back to Arabic. We, then, combined both the original dataset and the back-translated set.

We only back-translated the ArabGlossBERT training dataset (152,035 pairs). The testing dataset (15,172 pairs) was not back-translated, because it is used as an evaluation benchmark to compare

the performance improvement between the original and augmented datasets.

Table 2 provides statistics about the original ArabGlossBERT dataset, the newly added back-translations, and the whole dataset after augmentation. The augmentation shows that the size of the original dataset was doubled as it contains the original context-gloss pairs and the translated context-gloss pairs (152,032).

The original training dataset is not balanced with 55,585 positive pairs (36.6%) and 96,450 negative pairs (63.4%). To produce a more balanced dataset, we generated an additional 32,839 positive pairs by matching the original glosses with the new back-translated glosses increasing the number of positive pairs to 144,009. The 144,009 include 55,585 pairs from the original data, 55,585 pairs from back-translation and the added 32,839 pairs, resulting in a new dataset with 42.7% positive and 57.3% negative pairs.

Observations on Google Translate: First, although the quality of Google translations was generally acceptable, there are wrong translations. However, we did not make any improvements or revisions to these translations, as the goal of this paper is to measure whether automated back-translations can improve the accuracy of the trained models. Second, the output of the Google translation API was not always complete. In some cases it translates part of the input sentence. To overcome this challenge we used two techniques: 1) add special characters (.#) at the end of each sentence, if the special characters were translated back, then we know the translation reached the end of the sentence, 2) compare the length of the original and back-translated sentences and if the difference is significant, then this is an indication of incomplete translations. Partial translations are reviewed manually.

	Original ArabGlossBERT	Back-Translation Pairs	Augmented ArabGlossBERT
Unique un-diacritized lemmas	26,169	–	26,169
Unique glosses	32,839	32,839	65,678
Unique contexts	60,272	60,272	120,544
Training pairs	152,035	152,035 + 32,839	336,909
Positive pairs	55,585	55,585 + 32,839	144,009
Negative pairs	96,450	96,450	192,900
Testing pairs	15,172	–	15,172
Positive pairs	4,738	–	4,738
Negative pairs	10,434	–	10,434
Total: Train+Test	167,207	152,035 + 32,839	352,081

Table 2: Statistics of the original and augmented datasets.

Dataset	Description	Positive Pairs	Negative Pairs	Total
D_1	The original ArabGlossBERT dataset	55,585	96,450	152,035
D_2	D_1 with target signal	55,585	96,450	152,035
D_3	D_1 with context replaced by back-translated context	55,585	96,450	152,035
D_4	D_1 + Positive pairs of D_3	111,170	96,450	207,620
D_5	D_1 + D_3	111,170	192,900	304,070
D_6	D_1 + Positive pairs (original gloss - back-translated gloss)	88,424	96,450	184,874
D_7	D_4 + Positive pairs (original gloss - back-translated gloss)	144,009	96,450	240,459
D_8	D_5 + Positive pairs (original gloss - back-translated gloss)	144,009	192,900	336,909
D_9	D_1 + Positive pairs (original context - back-translated gloss)	111,170	96,450	207,620
D_{10}	D_1 + Pairs of cross relating the glosses against each other	88,424	373,955	462,379
D_{11}	D_1 (excluded pairs of functional words)	54,878	92,730	147,608
D_{12}	D_1 (only the pairs of the noun POS)	36,487	37,998	74,485
D_{13}	D_1 (only the pairs of the verb POS)	18,178	54,945	73,123

Table 3: The datasets that were used for the different experiments to fine-tune AraBERT on the TSV task.

4 Experiments

This section presents 13 experiments to measure the impact of data augmentation using back-translation on TSV model accuracy. The first experiment uses the original ArabGlossBERT dataset, D_1 , (as a baseline), while the other experiments are conducted with different dataset configurations. In all experiments, we used the original test dataset 15,172 pairs (4,738 positive and 10,434 negative). Table 3 presents the training datasets that we used in the experiments.

In all experiments we fine-tuned AraBERTv2 (aubmindlab/bert-base-arabertv02, CC-BY-SA) using the following hyperparameters: $\eta = 2e^{-5}$, batch size $B = 16$, max sequence length of 512, warm-up steps 1,412 and number of epochs 4.

The results of the 13 experiments are presented in Table 4, which includes precision, recall, F1-

score, and accuracy. The results are presented at the POS tag level and overall. Also, note that the test dataset is the same test set used in the original ArabGlossBERT dataset because we consider ArabGlossBERT as a baseline. In the next sub-sections, we elaborate on each experiment.

4.1 Experiment 1: D_1 Dataset (Baseline)

This experiment is the baseline for results comparison. We used the original dataset ArabGlossBERT, D_1 , without any augmentation and achieved the same results (83% accuracy) as reported in (Al-Hajj and Jarrar, 2022). Additionally, we evaluated the model performance per POS tag since the tokens are annotated with the POS tags (noun, verb, and functional words). While the accuracy across all tags is very similar (Table 4.), we observe a big difference in the Positive pair F1-score. For the

Dataset	Metric	All POS		Accuracy	Noun		Accuracy	Verb		Accuracy	Functional Words		Accuracy
		Positive	Negative		Positive	Negative		Positive	Negative		Positive	Negative	
D1 Baseline 152,035 pairs	Precision	76	85	83	75	85	82	78	85	83	63	84	81
	Recall	66	90		70	88		65	91		46	92	
	F1-Score	71	88		72	82		71	88		53	88	
D2 152,035 pairs	Precision	81	85	84	79	85	83	82	85	84	71	82	81
	Recall	65	93		68	91		64	94		36	95	
	F1-Score	72	89		73	88		72	89		48	88	
D3 152,035 pairs	Precision	68	80	77	65	79	75	70	80	78	55	79	77
	Recall	52	88		54	85		52	90		19	95	
	F1-Score	59	84		59	82		60	85		29	86	
D4 207,620 pairs	Precision	80	81	81	79	80	80	81	81	81	69	80	79
	Recall	53	94		55	92		53	94		23	97	
	F1-Score	64	87		65	86		64	87		34	88	
D5 304,070 pairs	Precision	76	82	81	77	79	80	76	84	82	70	80	79
	Recall	57	92		53	92		62	91		24	97	
	F1-Score	65	87		63	85		68	87		36	88	
D6 184,874 pairs	Precision	76	85	83	76	84	81	76	87	84	71	82	81
	Recall	67	90		66	89		70	90		32	96	
	F1-Score	71	88		71	86		73	88		44	88	
D7 240,459 pairs	Precision	79	82	81	77	81	80	80	83	82	71	79	79
	Recall	56	93		57	91		58	93		17	98	
	F1-Score	66	87		66	86		67	88		17	98	
D8 336,909 pairs	Precision	80	81	81	79	80	80	81	81	81	69	80	79
	Recall	54	94		55	92		53	94		23	97	
	F1-Score	65	87		65	86		64	87		34	88	
D9 207,620 pairs	Precision	78	84	83	77	83	81	78	86	84	73	81	81
	Recall	63	92		62	91		66	92		31	96	
	F1-Score	70	88		69	86		72	88		43	88	
D10 462,379 pairs	Precision	71	80	78	70	78	76	71	81	79	66	79	78
	Recall	51	90		50	89		54	90		19	97	
	F1-Score	59	85		58	83		61	85		30	87	
D11 147,750 pairs	Precision	80	81	81	79	80	80	81	81	81			
	Recall	54	94		55	92		53	94				
	F1-Score	65	87		65	86		64	87				
D12 74,485 pairs	Precision				80	82	81						
	Recall				60	92							
	F1-Score				69	87							
D13 73,123 pairs	Precision							74	84	81			
	Recall				62	90							
	F1-Score				68	87							

Table 4: Results, expressed as percentage, of the experiments for fine-tuning AraBERT on different combinations of the original ArabGlossBERT and augmented datasets.

functional words, the F1-score for Positive pairs is only 53%, compared to 72% and 71% for the nouns and verbs, respectively. We will notice this trend across all experiments, since functional words are highly polysemous (e.g., the preposition (في / in) has ten different glosses), and their glosses represent function and use in the sentence, rather than semantics.

4.2 Experiment 2: D_2 Dataset

The idea of this experiment is to use a learning signal by marking up the target word, in its context, with an open-close tag (<token>Target</token>) to emphasize the model learning of the target word. Thus, the dataset D_2 is the same as D_1 but with a learning signal surrounding the target words. This experiment is the same experiment conducted in (Al-Hajj and Jarrar, 2022) and we achieved the same results (84% accuracy). Overall, we see a 1% increase by using D_2 over D_1 . We note that D_2 is the only dataset with the target signal added.

4.3 Experiment 3: D_3 Dataset

This experiment evaluates the model performance using D_3 , which contains the back-translated context and the original gloss pairs (152,035). As shown in Table 4, the overall accuracy creased from 83% on D_1 to 77% on D_3 . The 6% drop in the accuracy illustrates that the quality of the back-translations is acceptable as an augmentation to the original data.

4.4 Experiment 4: D_4 Dataset

D_4 is original dataset D_1 in addition to the 55,585 Positive back-translated pairs. The motivation of adding the Positive back-translated pairs is to balance the original dataset, D_1 . Recall that D_1 contains 55,585 Positive pairs (36.6%) and 96,450 Negative pairs (63.4%) and by adding the Positive back-translated pairs, D_4 size increases to 207,620 pairs, among which 111,170 (53.5%) are positive pairs. Table 4 shows that this data configuration did not improve the model performance. On the contrary, the accuracy dropped by 2% compared to D_1 (baseline). We also note that the F1-score dropped from 71% to 64% for Positive pairs, and from 88% to 87% for Negative pairs.

4.5 Experiment 5: D_5 Dataset

D_5 consists of the original dataset D_1 in addition to its back-translation dataset D_3 . Although D_5 is

large (304,070 pairs), its accuracy is 81%, which is 2% lower than the baseline.

4.6 Experiment 6: D_6 Dataset

The D_6 dataset used in this experiment contains the original dataset D_1 , in addition to 32,839 Positive pairs that we generated by paring an original gloss with its back-translation. We achieved the same accuracy as the baseline (83%), but we believe that the fine-tuned model on D_6 is a little better than the baseline model for two reasons. First, the the F1-score for *noun* Negative pairs increased by 4% compared to the baseline to 86%, and the F1-score for *verb* Positive pairs increased by 2% to 73%. Second, since the training dataset is larger it is assumed to cover more vocabulary.

4.7 Experiments 7-8: D_7 and D_8 Datasets

Although we increased the size of datasets in these two experiments, their model accuracy and F1-scores are very similar, but lower compared with the baseline. D_7 contains the original dataset, the Positive back-translated pairs and the Positive glosses with their back-translations. With this data, we increased the Positive pairs to be 144,009 (60%) of the dataset. In experiment 8 we used D_8 , which contains the original dataset, all back-translation pairs, and the Positive gloss-gloss pairs.

4.8 Experiment 9: D_9 Dataset

D_9 contains D_1 and the 55,585 Positive pairs that we produced by pairing the original context with their back-translated gloss. The Positive pairs in D_9 account for 53.5% of the dataset. This data configuration achieved the same as the baseline (83% accuracy). Although the performance is same as the baseline, we see similar behaviour and we conclude the same as we did on the dataset D_6 .

4.9 Experiment 10: D_{10} Dataset

In this experiment we did not use back-translation. However, we augmented the original dataset D_1 such that, the set of glosses of a certain lemma are cross-related and the resulting pairs are considered Negative pairs. In this way, we were able to generate 32,839 Positive pairs and 277,505 Negative pairs, a total of 310,344 pairs. We augmented these pairs with D_1 resulting in 462,379 pairs. Notice that this is the hardest dataset to model because some negative pairs are generated at the lemma level and are harder to distinguish from their positive counterparts. The idea is to fine-tune a model

to be more sensitive in distinguishing positive and negative pairs, which as expected resulted in the lowest performance (78% accuracy) compared to other models.

4.10 Experiment 11: D_{11} Dataset

The goal of this experiment is to fine-tune a model excluding all pairs that are labeled with functional words. Functional words such as (إذا، عن، على، في، إلى) play the role of particles rather than providing core semantics. Additionally, they are frequently used in contexts and are highly polysemous. We fine-tuned a model with the D_{11} dataset, which is the same as the original dataset D_1 , but it excludes 4,427 pairs of functional words. However, the performance did not improve compared to the baseline. This illustrates that keeping the pairs of functional words is better than excluding them.

4.11 Experiment 12-13: D_{12} and D_{13} Datasets

The goal of these two experiments is to evaluate the pairs labeled with *nouns* and *verbs* separately. D_{12} contains 74,485 pairs, in which targets are *nouns* only, and D_{13} contains 73,123 pairs with *verb* targets. We fine-tuned two separate models for each of the datasets and achieved similar accuracy and F1-scores, however, the performance is slightly lower compared to the baseline. Nevertheless, since both D_{12} and D_{13} achieved similar results, we believe that fine-tuning the model on data with both POS tags allows for cross learning and in turn yields better performance.

5 Discussion and Conclusions

We presented an approach to improve Arabic TSV using automatic back-translation. We augmented an existing Arabic TSV dataset, ArabGlossBERT, by doubling its size with back-translated data using Google Translate API. To measure the impact of the data augmentation, we presented 13 experiments with different data configurations. Although we did not outperform the overall performance of the baseline model, we did observe that some experiments such as D_6 outperformed the baseline on *noun* positive pair and *verb* negative pair classification. Overall, our results are close to the results presented in (Lin and Giambi, 2021), which used back-translation augmentation for English TSV and achieved only 2% F1-score improvement. Nevertheless, we would like to note the following findings:

- Fine-tuning a BERT model using only the back-translation pairs achieved 77% accuracy (experiment 3), which is only 6% less than the baseline accuracy. This illustrates that the quality of automatic translations of glosses and contexts is not high but is generally acceptable.
- The different augmentations to the original dataset achieved between 78% to 83% accuracy (see experiments 4-9), but it did not outperform the baseline model. At the same time, augmentation did not harm the performance since the results are comparable to the baseline. Nevertheless, experiments 6 and 9 have illustrated a small improvement in the F1-scores for *noun* and *verb* POS. In addition, because D_6 and D_9 are larger than the baseline D_1 dataset, the fine-tuned models are assumed to cover a larger vocabulary and more contexts.
- Looking at the F1-scores, we note that the Positive pairs are always lower than the Negative pairs in all experiments and for all POS categories. This means that all models are less accurate at predicting Positive pairs. Although we tried to augment the dataset by increasing the number of Positive pairs, the F1-scores did not improve.
- In our attempts to fine-tune different models for each POS category, we found that: (1) excluding the pairs of functional words from the dataset (experiment 11) did not improve the performance, and (2) fine-tuning a model for all POS categories allows for cross learning from different POS tagged targets and yields better performance than fine-tuning separate models for *nouns* and *verbs* (experiments 12-13).

6 Limitations and Future Works

Our data augmentation as well as the experiments are based on (1) the quality of Google Translate API, (2) the quality of the glosses and contexts in the ArabGlossBERT training dataset, and most importantly on (3) the quality and coverage of the ArabGlossBERT test dataset. Although the quality of machine translation is limited, the goal of this paper is to measure whether such limited translations can improve the accuracy of the TSV fine-tuned

models. Additionally, the quality of the glosses and contexts in the ArabGlossBERT training dataset cannot be improved since they originated from Arabic lexicons. However, we believe that enriching the ArabGlossBERT by collecting more pairs from Arabic lexicons (i.e., building a rich Arabic sense inventory) will empower research on TSV and WSD tasks. More importantly, all experiments conducted in this paper used the ArabGlossBERT test dataset. Since there are no other testing datasets or benchmarks, the evaluation of our fine-tuned models is limited based on the quality and coverage of the ArabGlossBERT test dataset.

Next, we plan to develop another test dataset to evaluate our models and their generalizability. We plan to further explore other approaches for WSD task such as ranking of glosses, rather than addressing the WSD task through TSV.

Acknowledgment

We would like to thank Taymaa Hammouda and Ala Omar for the technical support on many aspects of this research.

References

- Moustafa Al-Hajj and Mustafa Jarrar. 2021. [Lu-bzu at semeval-2021 task 2: Word2vec and lemma2vec performance in arabic word-in-context disambiguation](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 748–755, Online. Association for Computational Linguistics.
- Moustafa Al-Hajj and Mustafa Jarrar. 2022. [Arabglossbert: Fine-tuning bert on context-gloss pairs for wsd](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021), Held Online, 1-3September, 2021*, pages 35–43.
- 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.
- Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, Roberto Navigli, et al. 2021. [Recent trends in word sense disambiguation: A survey](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4330–4338.
- Jan A Botha, Zifei Shan, and Dan Gillick. 2020. [Entity linking in 100 languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7833–7845.
- Anna Breit, Artem Revenko, Kiamehr Rezaee, Mohammad Taher Pilehvar, and Jose Camacho-Collados. 2020. [Wic-tsv: An evaluation benchmark for target sense verification of words in context](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 1635–1645.
- Kareem Darwish, Nizar Habash, Mourad Abbas, Hend Al-Khalifa, Husein T. Al-Natsheh, Houda Bouamor, Karim Bouzoubaa, Violetta Cavalli-Sforza, Samhaa R. El-Beltagy, Wassim El-Hajj, Mustafa Jarrar, and Hamdy Mubarak. 2021. [A panoramic survey of natural language processing in the arab worlds](#). *Commun. ACM*, 64(4):72–81.
- Erik De Vries, Martijn Schoonvelde, and Gijs Schumacher. 2018. [No longer lost in translation: Evidence that google translate works for comparative bag-of-words text applications](#). *Political Analysis*, 26(4):417–430.
- 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.
- Mohammed El-Razzaz, Mohamed Waleed Fakh, and Fahima A Maghraby. 2021. [Arabic gloss wsd using bert](#). *Applied Sciences*, 11(6):2567.
- Bradley Hauer and Grzegorz Kondrak. 2022. [Wic=tsv= wsd: On the equivalence of three semantic tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2478–2486.
- Luyao Huang, Chi Sun, Xipeng Qiu, and Xuan-Jing Huang. 2019. [Glossbert: Bert for word sense disambiguation with gloss knowledge](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3509–3514.
- Mustafa Jarrar. 2006. [Towards the notion of gloss, and the adoption of linguistic resources in formal ontology engineering](#). In *Proceedings of the 15th international conference on World Wide Web (WWW2006)*, pages 497–503. ACM Press, New York, NY.
- Mustafa Jarrar. 2011. [Building a formal arabic ontology \(invited paper\)](#). In *Proceedings of the Experts Meeting on Arabic Ontologies and Semantic Networks*. ALECSO, Arab League.

- Mustafa Jarrar. 2018. [Search engine for arabic lexicons](#). In *Proceedings of the 5th Conference on Translation and the Problematics of Cross-cultural Understanding*. The Forum for Arab and International Relations.
- Mustafa Jarrar. 2020. [Digitization of Arabic Lexicons](#), pages 214–217. UAE Ministry of Culture and Youth.
- Mustafa Jarrar. 2021. [The arabic ontology - an arabic wordnet with ontologically clean content](#). *Applied Ontology Journal*, 16(1):1–26.
- Mustafa Jarrar and Hamzeh Amayreh. 2019. [An arabic-multilingual database with a lexicographic search engine](#). In *The 24th International Conference on Applications of Natural Language to Information Systems (NLDB 2019)*, volume 11608 of *LNCS*, pages 234–246. Springer.
- Mustafa Jarrar, Hamzeh Amayreh, and John P. McCrae. 2019. [Representing arabic lexicons in lemon - a preliminary study](#). In *The 2nd Conference on Language, Data and Knowledge (LDK 2019)*, volume 2402, pages 29–33. CEUR Workshop Proceedings.
- Mustafa Jarrar, Mohammed Khalilia, and Sana Ghanem. 2022. [Wojood: Nested arabic named entity corpus and recognition using bert](#). In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2022)*, Marseille, France.
- Mustafa Jarrar, Fadi Zaraket, Rami Asia, and Hamzeh Amayreh. 2018. [Diacritic-based matching of arabic words](#). *ACM Asian and Low-Resource Language Information Processing*, 18(2):10:1–10:21.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Guan-Ting Lin and Manuel Giambi. 2021. [Context-gloss augmentation for improving word sense disambiguation](#). *arXiv preprint arXiv:2110.07174*, abs/2110.07174.
- George A Miller, Claudia Leacock, Randee Teng, and Ross T Bunker. 1993. [A semantic concordance](#). In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Jose G Moreno, Elvys Linhares Pontes, and Gaël Dias. 2021. [Cttr@ wic-tsv: Target sense verification using marked inputs and pre-trained models](#). In *Proceedings of the 6th Workshop on Semantic Deep Learning (SemDeep-6)*, pages 1–6.
- Nikhil Patel, James Hale, Kanika Jindal, Apoorva Sharma, and Yichun Yu. 2021. [Building on huang et al. glossbert for word sense disambiguation](#). *arXiv preprint arXiv:2112.07089*, abs/2112.07089.
- Niloofer Ranjbar and Hossein Zeinali. 2021. [Lotus at semeval-2021 task 2: Combination of bert and paraphrasing for english word sense disambiguation](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 724–729.
- Boon Peng Yap, Andrew Koh, and Eng Siong Chng. 2020. [Adapting bert for word sense disambiguation with gloss selection objective and example sentences](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 41–46.