

# Wojood: Nested Arabic Named Entity Corpus and Recognition using BERT

Mustafa Jarrar, Mohammed Khalilia, Sana Ghanem

Birzeit University

Palestine

{mjarrar, mkhalilia, swghanem}@birzeit.edu

## Abstract

This paper presents Wojood, a corpus for Arabic *nested* Named Entity Recognition (NER). Nested entities occur when one entity mention is embedded inside another entity mention. Wojood consists of about 550K Modern Standard Arabic (MSA) and dialect tokens that are manually annotated with 21 entity types including person, organization, location, event and date. More importantly, the corpus is annotated with nested entities instead of the more common flat annotations. The data contains about 75K entities and 22.5% of which are nested. The inter-annotator evaluation of the corpus demonstrated a strong agreement with Cohen’s Kappa of 0.979 and an F1-score of 0.976. To validate our data, we used the corpus to train a nested NER model based on multi-task learning using the pre-trained AraBERT (Arabic BERT). The model achieved an overall micro F1-score of 0.884. Our corpus, the annotation guidelines, the source code and the pre-trained model are publicly available.

**Keywords:** Named Entity Recognition, Multi-Task Learning, Nested Entities, BERT, Arabic NER Corpus

## 1. Introduction

Named Entity Recognition (NER) is integral to many Natural Language Processing (NLP) applications such as chatbots and question answering (Shaheen and Ezzeldin, 2014), information retrieval (Guo et al., 2009), machine translation (Hassan and Sorensen, 2005), word-sense disambiguation (Al-Hajj and Jarrar, 2021a), clustering (Abuaiadah et al., 2017), interoperability (Jarrar et al., 2011), topic modeling and event discovery (Feng et al., 2018), among others. NER is the task of identifying named entity mentions in unstructured text and classify them to a predefined categories such as *person*, *organization*, *event*, *location*, or *date*. For example, in a sentence like "Sami works at Jimmy Carter Center", "Sami" would be tagged as a *person*, and "Jimmy Carter Center" as an *organization*, as shown in Figure 1. We use the IOB2 tagging scheme, where *B* indicates the beginning of the entity mention, *I* indicates the inside token, and *O* refers to outside entity (Sang and Veenstra, 1999). The NER task becomes a sequential labeling task, where each token is tagged with a label based on the contextual information surrounding it (Sohrab and Miwa, 2018).




Figure 1: Flat NER Example

Many approaches were proposed for NER in many languages, including Arabic, which we will review in Section 2. However, less attention is given to *nested named entities* (Finkel and Manning, 2009), which are entity mentions contained inside other entity mentions. For example, "Jimmy Carter" is a *person* that is inside the

"Jimmy Carter Center" *organization*, and "Birzeit" is a *geopolitical* entity inside the organization mention "Birzeit University". More examples in Arabic are shown in Figure 2.

There are a few reasons why nested NER is understudied. First and most important is lack of data, specially for low-resource languages like Arabic (Darwish et al., 2021; Jarrar, 2021). In English, there are a few datasets tagged with nested entities. The GENIA molecular biology corpus (Ohta et al., 2002) with 400K tokens and almost 100K annotations, but only 17% of the entities are embedded in another entity (Katiyar and Cardie, 2018). The ACE2005 which contains about 30K mentions among which there are 3,300 nested mentions (Walker et al., 2006; Lu and Roth, 2015). The NNE newswire corpus with 279K mentions, 114 entity types with up to six layers of nesting (Ringland et al., 2020). However, in Arabic language, which is the focus of this paper, there are a handful of datasets annotated with flat entities, but none are annotated with nested entities. Second, nested NER adds complexity to the machine learning models in terms of computation, runtime and accuracy (Finkel and Manning, 2009; Sohrab and Miwa, 2018). Such models need to recursively predict the next tag. Third, data annotation is very expensive and time consuming. An effective nested NER dataset should cover multiple domains, be large in size and have multiple layers of nesting. Also, nested entity annotations are more prone to human errors than flat annotations, since there are multiple layers, which means we need to invest more in inter-annotator agreement to ensure high-quality data.

The goal of this paper is to provide some solutions to overcome some of the challenges we mentioned earlier.

Our main contribution is *Wojood*<sup>1</sup>, a relatively large-scale Arabic nested named entity corpus. The corpus consists of 550K tokens, which were manually annotated. The corpus supports 21 entity types, four layers of nesting and covers both MSA and dialect in multiple domains (see statistics in Table 2). Our evaluation of the inter-annotator agreement shows an overall *Kappa* of 0.979 and an F1-score of 0.976. Furthermore, the paper presents a method for fine-tuning AraBERT (Arabic BERT) on a multi-task learning architecture to predict nested tags. The model achieved an overall micro F1-score of 0.884.

The rest of this article is organized as follows: Section 2 presents related work, Section 3 presents our proposed corpus, Section 4 describes the BERT-based multi-task architecture and results, Section 5 gives overview of the implementation of our NER web service, and we conclude in Section 6.

## 2. Related Work

There are a few available Arabic NER corpora (summarized in Table 1). However, none of them supports nested entities and most of them are small in size and support only a small number of entity types. The richest corpus is Ontonotes 5, which is available through the LDC catalog with restricted access (Weischedel et al., 2013). It consists of  $\sim 300$ K tokens, collected from MSA media sources and annotated with 18 entity types. The ANERCorp corpus was also collected from MSA media sources (Benajiba et al., 2007). It contains  $\sim 150$ K tokens and supports only four entity types (person, organization, location, and miscellaneous). The Classical Arabic Named Entity Recognition Corpus (CANERCorpus) is a Classical Arabic of Hadith documents ( $\sim 258$ K tokens), which was annotated manually using 14 entity types; however, most entity types are religion-specific (Salah and Zakaria, 2018). Although the corpus contains about 72K entity mentions, about 93% of these entities are tagged as Allah, Prophet, persons, and numbers. The AQMAR corpus was collected from Wikipedia and covers four domains: history, technology, science, and sports (Mohit et al., 2012). The corpus consists of  $\sim 47$ K tokens and was annotated with open-ended types of entities. Furthermore, Mostefa et al. (2009) presented a multilingual NER corpus that was collected from news articles and included 56 million Arabic tokens. The corpus was, first, automatically annotated with five entity types using a rule-based approach, then a subset of the annotations was verified manually.

Compared with these corpora, *Wojood* is larger, supports nested entities, 21 entity types, and multiple domains (health, finance, technology, law, elections, politics, migration, terrorism, history, culture, and social media topics). Our corpus also includes MSA and dialectal text.

Corpus	Nested	Tokens	Entities	Classes	Arabic	Domain
Ontonotes5	No	300k	28k	18	MSA	News
ANERCorp	No	150k	11k	4	MSA	News
CanerCorpus	No	258k	72k	14	Classic	Religion
AQMAR	No	74k	-	open	MSA	4 domains
<b>Wojood Corpus</b>	<b>YES</b>	<b>550K</b>	<b>75K</b>	<b>21</b>	<b>MSA-Dialect</b>	<b>Multi</b>

Table 1: Existing Arabic NER Corpora

For entity recognition, several NER approaches have been proposed for Arabic as reviewed in (Shaalán, 2014; Ali et al., 2020). NER approaches in Arabic can be grouped into rule-based, machine learning, and advanced deep learning approaches. Rule-based approaches use linguistic features (context, morphology, character-level, etc.) and maybe gazetteers for detecting entities (Shaalán and Raza, 2007; Jaber and Zaraket, 2018). Machine Learning methods such as support vector machines (SVM) and conditional random field (CRF) were also used for NER with less to no need for language-specific features (Benajiba et al., 2007; Abdul-hamid and Darwish, 2010; Mohit et al., 2012; Abdallah et al., 2012). More recently, Deep Learning approaches have shown better results in recognizing entities. For example, an F1-score of 0.8877 was achieved with the ANERcorp corpus using convolution neural networks (CNN) and bi-directional long-short term memory (BiLSTM) (Khalifa and Shaalan, 2019). Experiments with BERT models, using the same ANERcorp corpus, achieved 0.8438 and 0.8913 F1-score with AraBERT2 (Antoun et al., 2020) and ARBERT (Abdul-Mageed et al., 2021), respectively. A CRF layer was added to AraBERT2 by (Al-Qurishi and Souissi, 2021), which increased the F1-Score to 0.912. Nevertheless, none of these approaches tackled the problem of recognizing nested Arabic named entities due to the absence of such corpora for Arabic.

NER has achieved a remarkable accuracy on many datasets in high-resource languages such as English. Models were published on CoNLL-2002 (Tjong Kim Sang, 2002), CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003), Ontonotes (Weischedel et al., 2011) and i2b2 clinical dataset (Uzuner et al., 2011). CoNLL has approximately 320K tokens annotated with four entity types: person, location, organization and miscellaneous (Passos et al., 2014). Ontonotes contains about 1.6M tokens annotated with 18 entity types (Passos et al., 2014) and 2010 i2b2 dataset contains about 416K tokens (Dai and Adel, 2020).

One of the most common model architectures for NER models is based on LSTM and CRF (Lample et al., 2016). LSTM based architecture was extended by others for clinical joint entity extraction and assertion detection such as negation (Bhatia et al., 2019). Biomedical NER based on Recurrent Neural Network (RNN) was proposed in (Li et al., 2015). LSTM and CNN

<sup>1</sup> *Wojood* Corpus: <https://ontology.birzeit.edu/wojood>

based models rely on pre-trained word embeddings such as Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014) and FastText (Bojanowski et al., 2017).

More recently, (Devlin et al., 2019) proposed a new language representation model, bi-directional representation from transformers (BERT). Transformer based language models (LM) rely on attention to capture the contextual information and they can model longer dependencies. A common BERT model for NER usually consists of the BERT encoder to encode each token in the input sequence into a 768 dimensional vectors, which are used as input to a linear layer, followed by an activation function such as softmax to output a probability distribution of the tags for each token.

Although nested NER is more complex, it started to receive greater attention in recent years in languages like English, thanks for the advancements in machine learning and deep learning. A multi-layer BiLSTM that learns hypergraph representation for nested entities proposed in (Katiyar and Cardie, 2018), which was trained on ACE2004 and ACE2005 datasets and the authors reported 0.727 and 0.705 F1-score, respectively. A stacked flat LSTM model proposed by (Ju et al., 2018) for nested tag prediction. The model consists of an embedding layer, followed by stacked flat NER units, each unit contains an LSTM and CRF layer. The number of stacked layers is dynamically added until the model predicts "O" tags for all tokens in the sequence. The output of the current flat NER layer are merged to form a representation of the detected entities, which are fed into the next flat NER layer. The model assumes the sequence is tagged once with the same entity type. The author reported F1-score of 0.771 on the nested GENIA molecular biology corpus (Ohta et al., 2002). A similar approach to (Ju et al., 2018) was proposed by (Wang et al., 2020) uses a pyramid architecture. The sequence is recursively encoded through stacked flat NER layers from bottom to top. Similarly, the model relied on LSTM architecture to encode the sequence. Their model was trained on ACE2004, ACE2005, GENIA, and NNE newswire, with 0.8774, 0.8634, 0.7931, and 0.9468 F1-scores, respectively.

Another LSTM based model published by (Sohrab and Miwa, 2018) uses an exhaustive approach. The model exhaustively considers all possible regions in the sequence, however, the maximum number of regions is user pre-defined. Their approach allows for predicting nested tag sequences for entities of the same type. The authors in (Shibuya and Hovy, 2020) proposed a nested NER LSTM-CRF model that uses second-best sequence strategy to predict nested tags. In their approach, the LSTM layer acts as the encoder, followed by CRF for sequence decoding. Their approach uses the probabilities returned by CRF to compute the next best sequence of tags. The list above is certainly not exhaustive, but it illustrated the need for nested NER datasets and models in the Arabic language.

### 3. The Corpus

#### 3.1. Corpus Collection and Preparation

Our corpus was collected from multiple sources, various domains and covers a number of different topics (Table 2). Subset of the corpus (~258K tokens) was manually collected from online articles published on websites including un.org, hrw.org, msf.org, who.org, mipa.institute, elections.ps, sa.usembassy.gov, diplomatie.ma, and quora.com. This subset covers several topics including health, information and communication technologies (ICT), finance, elections, law, politics, migration, and terrorism. The second subset (~227K tokens) was extracted from Awaq, the Birzeit University digital Palestinian archive, which covers modern history and cultural heritage. We extracted the descriptions of about 17K documents. Each description is one sentence (4-30 tokens) describing the content of a document. The importance of this subset is that it covers many types of entities and about 60% of the tokens are entities, most of which are overlapping. The third subset (~65K tokens) is based on the Palestinian (Jarrar et al., 2017; Jarrar et al., 2014) and Lebanese (Haff et al., 2022) corpora, a dialect text collected from social media and other Levant resources discussing general topics. We aimed at generating a corpus that contains both dialect text and MSA for diversification.

No cleaning, normalization, or pre-processing steps were performed. The raw text was segmented into sentences, each is given a sentence ID. Each sentence was then tokenized, and each token is assigned an index to represent its position in the sentence. The corpus was then saved using a comma separated values (CSV) format and uploaded to Google Sheets for annotation. The Sheet is a table of four columns: sentence ID, token position, token, and the NER tags. The annotators were asked to read the sentence vertically, and tag each token in the fourth column. Table 5 shows a tokenized sentence and its tags.

Source - Topics	Sentences	Tokens
Web Articles (MSA) Health, Finance, ICT, Law, Elections, Politics, Migration and Terrorism	9,053	258,102
Archive (MSA) History and Culture	12,271	227,020
Social Media (Dialect) General topics	5,653	65,342
<b>Total</b>	<b>26,977</b>	<b>550,464</b>

Table 2: Statistics about the raw corpus

#### 3.2. Annotation Process

Twelve annotators were carefully selected to participate in the annotation process. All annotators are 4<sup>th</sup>-year students with high distinction at Birzeit University. A half-day workshop was organized to train




Figure 2: Example of nested annotations. (a) nested entity of mentions of different types (“Bank of Cairo Manager was awarded one million pound”), (b) illustrates nested entity mentions with same type (“Employees’ union at Cairo University to support 2021 budget”).

about 30 candidate annotators and explain the annotation guidelines; then, each was given a quiz of 750 tokens to annotate, which we designed to be hard. The 12 students who achieved above 90% in the quiz were recruited. The annotation process, managed by two NER experts, was conducted in three phases over eight months:

**Phase I:** Each annotator was given about 46K tokens to annotate. The two experts were also helping on critical cases and providing continuous feedback to the annotators during this phase.

**Phase II:** The experts reviewed all annotations in phase I and asked the annotators to re-visit some annotations. A program was developed to automatically extract all entities from the annotated corpus in Phase I. The experts then carefully reviewed all extracted entities. In case of errors, doubts, or incomplete annotations, feedback is provided by the experts to the annotators. The goal here is to improve the consistency of the annotations across all annotators. In other words, our main goal in this phase is not to detect missing annotations; but rather, to validate the entities that have been annotated and increase the alignment among the annotators.

**Phase III:** We fine-tuned BERT on a multi-task learning architecture to predict nested entities based on the annotated corpus then used this model to re-annotate the corpus. We compared the manual annotations with the annotations suggested by the model and asked the annotators to verify the mismatches. We repeated this step twice. With this strategy we were able to detect some missing annotations.

Table 3 provides some statistics about the final annotations and the frequency of each entity type. The final

dataset contains about 75K entities, out of which 17K entities (22.5%) are nested within other entity mentions and only 576 of the nested entities are of the same type. Among all entity types, ORG, OCC and FAC are the entities that overlap the most.

Tag	Count Flat	Count Nested	Total
PERS	6531	739	7,270
NORP	4,928	334	5,262
OCC	5,351	164	5,515
ORG	15,292	3,493	18,785
GPE	11,501	10,279	21,780
LOC	755	162	917
FAC	939	276	1,215
PRODUCT	54	1	55
EVENT	2649	123	2,772
DATE	2,398	105	2,503
TIME	331	2	333
LANGUAGE	197	1	198
WEBSITE	607	0	607
LAW	496	0	496
CARDINAL	1,790	23	1,813
ORDINAL	4,041	989	5,030
PERCENT	137	0	137
QUANTITY	49	8	57
UNIT	5	54	59
MONEY	197	30	227
CURR	25	216	241
<b>Total</b>	<b>58,273</b>	<b>16,999</b>	<b>75,272</b>

Table 3: Counts of the flat, nested, and total of each entity type in the corpus.

### 3.3. Annotation Guidelines

The data was annotated with 21 entity types, see Table 4 for a short description and examples for each entity type. The full documentation of the guidelines is available online<sup>2</sup>. Our annotation guidelines are designed to be compatible with the guidelines proposed by the LDC’s OntoNotes 5 guidelines (Weischedel et al., 2013). However, we introduced four new tags (occupation, website, unit, currency), and revised the LDC guidelines to be more suitable for Arabic names. For example, we consider prefixes and suffixes to be part of entity names, such as (سادمهم-لخان الخليلي-محمد).

**Nested annotations** are the core of our corpus. If an entity mention is part of another entity mention, both are annotated. Figure 2 illustrates examples of nested annotations. For tokens that are annotated with multiple labels, these labels are sorted from outside to inside or top to bottom. For example, in Figure 2a the token (بنك /bank) is labeled first with I-OCC then B-ORG because the OCC is the top-most entity; thus, (B-ORG I-OCC) would be incorrectly ordered.

In some cases, which are rare, two entities of the same type may overlap, as illustrated in Figure 2b, which

<sup>2</sup>our guidelines: <https://ontology.birzeit.edu/wojood>

Tag	Short Description
PERS	People names, including first, middle, last and nicknames. Titles are not included except Prophets, kings, etc. فيروز، عادل إمام، أم كلثوم، ابن أحمد، الملك عبدالله، النبي محمد
NORP	Group of people. العرب، المسيحيين، سكان القدس، مصابي كورونا، وزراء الخارجية العرب
OCC	Occupation or professional title. الدكتور، المهندس، رئيس جامعة الزيتونة، مدير بنك القاهرة، قائد الجيش
ORG	Legal or social bodies like institutions, companies, agencies, teams, parties, armies, and governments. بنك القاهر، شركة أرمكس، ريال مدريد، داعش، الجيش المصري، حكومة العراق
GPE	Geopolitical like countries, cities, and states. ليبيا، صفاقس، مدينة القدس، الجمهورية اللبنانية، الولايات المتحدة الأمريكية
LOC	Geographical locations (Non-GPE), rivers, seas, mountains, and geographical regions. البحر الميت، قناة السويس، الخليج العربي، آسيا، الوطن العربي، شرق آسيا
FAC	Name of a specific place, like roads, cafes, buildings, airports, and gates. مطار صنعاء، حنين أبو غريب، المسجد الأقصى، شمال الحرة، مقهى الفيضاي
PRODUCT	Vehicles, weapons, foods, etc. مرسيدس بي 180، أيفون 13، إم 16، دبابة مراكيا، تروفين، ونستون
EVENT	Name of an event of general interest, like battles, wars, sports events, demonstrations, disasters, elections, and national or religious days. The place and date of the event are included as part of the event name. حرب عام 1973، القمة العربية لعام 2005، عيد الفطر، يوم الأرض
DATE	Reference to specific or relative dates including age, era, duration, month and day names. Charterers used to separate date components are part of the date. 13، 2022، يونيو، الاسبوع الماضي، من 2019 حتى 2020، الفترة العشانية
TIME	Reference to specific or relatives times, which is less than a day; including day times like evening and sunset. الساعة 21، من الساعة الخامسة حتى السابعة مساءً، خلال ساعتين، صباح الجمعة
LANGUAGE	Named human language or named dialect. اللغة العربية، الفصحى، الدارجة المغربية، الإنجليزية، اللغة الفرنسية
WEBSITE	Any named website or specific URL. صفحة جامعة الزيتونة، موقع فيسبوك، يوتيوب، www.schema.org
LAW	Reference to a specific law or part of it. المادة (114) من قانون العقوبات لسنة 2005، قانون الاستعمار، الدستور الأردني
CARDINAL	Numerals written in digits or words. 150، 30، 0.1، صفر، اثنان، أربعة وعشرون، مليون
ORDINAL	Any ordinal number, in digits or words, that does not refer to a quantity. 125، الخامس، الأول، ثانياً، ثالثاً، الرابع والعشرين، A2341
PERCENT	A word or a symbol referring to a percent. 5 بالمئة، 10 ٪، 9 من كل ألف
QUANTITY	Any value measured by standardized units, except dates, times, and money. 3 كيلومتر، مئة قدم، 3 طن، 50 غرام، 25 سم مربع
UNIT	Any name or symbol of a standardized unit. ميل، كيلو، كيلومتر، إنش، كيلوغرام، هكتار، مل، بيكو
MONEY	Absolute monetary quantity, including currency names. مئة وخمسون درهماً إماراتي، اثنان وثلاثون يورو، 8 دولار، 6€
CURR	Any name or symbol referring to currency. دولار، جنيه مصري، دينار، فرنك، ريال سعودي، S، €

Table 4: NER classes

Token	1 <sup>st</sup> Layer	2 <sup>nd</sup> Layer	3 <sup>rd</sup> Layer
لائتلاف /Union of	B-ORG		
العاملين /employees	I-ORG		
بجامعة /at University	I-ORG	B-ORG	
القاهرة /of Cairo	I-ORG	I-ORG	B-GPE
لدعم /to support	O		
ميزانية /budget	O		
2021	B-DATE		

Table 5: Examples of nested entity annotations

shows that (لائتلاف العاملين بجماعة القاهرة) /The employee union of Cairo University) is an organization that overlaps with (جامعة القاهرة) /Cairo University) which is another organization. When annotating nested entities of the same type, some token may have duplicate tags. Table 5 shows a sentence with three layers of nesting. Notice that (القاهرة) /Cairo) is part of three mentions, two of which are ORG. (القاهرة) /Cairo), which happened to be the last token in the entity mention is thus assigned I-ORG in both the 1<sup>st</sup> and 2<sup>nd</sup> layers.

### 3.4. Inter-annotator Agreement

To evaluate the Inter-Annotator Agreement (IAA), we asked each annotator to annotate 2,000 tokens that were previously annotated by a different annotator, and which we selected randomly, from different sources and domains. We also paid attention to select sentences from all domains. In this way, the IAA is performed between:  $A_1 - A_2$ ,  $A_2 - A_3$ , ...,  $A_{12} - A_1$ , with a total of about 24K token (4.3% of the corpus). Table 6 presents IAA for each entity type.

Although IAA is typically measured using Cohen’s *Kappa*, some researchers did not recommend it for evaluating named entity annotations (Hripcsak and Rothschild, 2005; Campillos-Llanos et al., 2021). They noted that named entities are sequences of tokens and do not contain negative cases. As stated in (Hripcsak and Rothschild, 2005), “*Kappa statistics and reliability coefficients cannot be calculated in studies without a negative case count*”. Additionally, since the corpus might be unbalanced, as the number of tokens tagged “O” is much higher than other labels, *Kappa* would be high and overestimates IAA. On the other hand, ignoring the “O” label yields low *Kappa* scores (Hripcsak and Rothschild, 2005; Brandsen et al., 2020; Campillos-Llanos et al., 2021; Martínez-deMiguel et al., 2022). Instead, they recommended to use the F1-score, as it might be more appropriate in reflecting the IAA in NER tasks. Nevertheless, in this paper, we present the three evaluations: (i) *Kappa* with “O”, denoted as  $\kappa_O$ , (ii) *Kappa* without “O”, denoted as  $\kappa_{\sim O}$ , and (iii) F1-score. In fact, we found that the three evaluation scores are close to each other.

### Calculating $\kappa_O$

To calculate the  $\kappa_O$  for a certain tag (e.g., ORG), we count the number of agreements and disagreements be-

tween annotators on this tag. All other tags are ignored and counted as "O". Agreements are counted as pairwise matches at the token level; thus, if a token is annotated, e.g., as ORG by one annotator and as O by another annotator, we count it as a disagreement. As such, *Kappa* is calculated as the following (Di Eugenio and Glass, 2004):

$$\kappa = \frac{P_o - P_e}{1 - P_e}$$

where  $P_o$  is the observed agreement between annotators, and  $P_e$  is the expected agreement (agreement by chance) defined as the agreement between annotators obtained if they randomly assign tags while annotating.  $P_e$  is calculated as:

$$P_e = \frac{1}{N^2} \sum_T n_{T1} \times n_{T2}$$

where  $n_{T_i}$  is the number of tokens to which annotator  $i$  assigned the tag  $T$ .  $N$  is the total number of annotated tokens. The overall  $\kappa_O$  is calculated using macro-average of all *Kappa* scores of the other tags.

### Calculating $\kappa_{\sim O}$

The  $\kappa_{\sim O}$  for a given tag is calculated similar to the calculations of the  $\kappa_O$ , however, we did not include the agreement on the "O" label.

### Calculating F1-score

To calculate the F1-score for a specific tag  $T$  (e.g., PERS), we counted only the tokens labeled with  $T$  by at least one of the annotators. Then, we performed a pair-wise comparison and counted the matches of this tag as true positives  $TP$  (i.e. number of tokens labeled PERS by both annotators). Otherwise, if the first annotator disagrees with the second, we count those as false negatives  $FN$ , and if the second disagrees with the first, we count them as false positives  $FP$ . In this way the number of disagreements is  $FN + FP$ . Calculating the F1-score for a certain tag  $T$  is then given by:

$$F1 - Score = \frac{2TP}{2TP + FN + FP}$$

The overall F1-score of all annotations is the micro-average of the F1-scores of all tags, which takes into account the count (i.e., weight) of each tag.

### 3.5. Annotation Challenges

The results in Table 6 illustrate a very strong agreement on all tags. Closer examination of the disagreement scores shows that EVENT, NORP, TIME, LOC, FAC and DATE have the most disagreements. We found that this is mostly due to entity boundaries. Although the annotation guidelines stated that some tokens, like (بداية /beginning of)

Tag	TP	FN	FP	$\kappa_O$	$\kappa_{\sim O}$	F1-Score
PERS	270	2	1	0.994	0.994	0.994
NORP	659	29	26	0.959	0.955	0.96
OCC	486	11	2	0.987	0.986	0.987
ORG	1,713	33	30	0.981	0.974	0.982
GPE	778	7	13	0.987	0.985	0.987
LOC	135	7	4	0.961	0.96	0.961
FAC	48	0	3	0.97	0.969	0.97
PRODUCT	5	0	0	1	1	1
EVENT	386	56	14	0.915	0.91	0.917
DATE	688	28	8	0.974	0.971	0.975
TIME	63	8	3	0.919	0.919	0.92
LANGUAGE	-	-	-	-	-	-
WEBSITE	-	-	-	-	-	-
LAW	257	1	0	0.998	0.998	0.998
CARDINAL	250	3	6	0.982	0.982	0.982
ORDINAL	277	1	4	0.991	0.991	0.991
PERCENT	43	0	0	1	1	1
QUANTITY	6	0	0	1	1	1
UNIT	3	0	0	1	1	1
MONEY	29	0	0	1	1	1
CURR	14	0	0	1	1	1
<b>Overall</b>	<b>6,110</b> count	<b>114</b> count	<b>186</b> count	<b>0.98</b> macro	<b>0.979</b> macro	<b>0.976</b> micro

Table 6: Overall inter-annotator agreement for each entity type.

e.g., (بداية /beginning of 2021), but some annotators excluded them. Similarly, (العرب المهاجرين /Arab emigrants) was annotated as one NORP entity by some annotators, while others annotated (العرب /Arab) and (المهاجرين /emigrants) as two separate NOPR entities. Most of the disagreements on ORG were related to the boundary issue as well. For example, the (OECD منظمة التنمية والتعاون الاقتصادي / Organisation for Economic Co-operation and Development), was annotated as one ORG entity by one annotator, while the (OECD) was annotated as a separate ORG entity by another annotator. The (البنك العربي بعمان /The Arab Bank in Amman) was sometimes annotated as an ORG entity, while others excluded (بعمان /Amman) from the entity mention.

The EVENT was problematic. Not only that annotators disagree on the boundaries, but they also disagree on the relevancy of the events. For example, one annotator annotated (اعترفت الولايات المتحدة بالملكية المتوكلية اليمنية عام ١٩٤٦ / The USA recognized the Mutawakkilite Kingdom of Yemen in 1946) as an EVENT, while others ignored the event completely. Indeed, some events might be important to annotate in some domains but not others.

We believe that the reason of why the overall inter-annotator agreement in our corpus was high is because we managed to provide continuous feedback to the annotators and revise the annotation guidelines if needed. Additionally, the review we carried out in Phase II was critical to improve the consistency among the annotators. It was also beneficial to compare the annotators and the model annotations, specially in the cases where annotators missed, but the model did not.

## 4. Named Entity Recognition (NER)

In addition to the inter-annotator agreement, we developed a model for nested NER to further validated the data and showcase its accuracy.

Tag	F1-score											
	1-2	2-3	3-4	4-5	5-6	6-7	6-8	8-9	9-10	10-11	11-12	12-1
PERS	1	1	0.97	1	-	1	1	1	1	0.95	1	-
NORP	0.92	0.97	0.92	0.93	0.96	0.94	100	0.97	0.95	0.96	1	1
OCC	0.98	1	0.97	0.92	0.98	1	1	0.98	0.99	1	1	1
ORG	0.96	0.99	0.94	0.95	0.98	0.98	1	1	0.98	0.99	1	1
GPE	1	1	0.98	0.98	0.98	0.96	1	0.99	0.95	1	1	1
LOC	1	1	0.92	1	-	0.91	1	1	0.98	1	-	-
FAC	1	-	1	1	-	1	1	1	1	-	-	-
PRODUCT	-	-	-	-	-	1	-	-	-	1	-	-
EVENT	-	1	0.90	0.87	0.92	0.91	1	0.92	1	1	-	1
DATE	1	1	0.97	0.97	1	0.92	1	0.94	0.98	0.92	1	1
TIME	-	1	1	1	-	-	1	1	-	-	1	-
LANGUAGE	-	-	-	-	-	-	-	-	-	-	-	-
WEBSITE	-	-	-	-	-	-	-	-	-	-	-	-
LAW	-	-	-	-	1	-	-	-	-	-	0.99	1
CARDINAL	0.98	1	0.97	0.99	1	1	1	93	1	0.96	1	1
ORDINAL	0.92	1	0.97	1	1	0.95	1	0.96	1	1	1	1
PERCENT	1	1	1	1	1	1	-	-	1	-	-	1
QUANTITY	-	-	-	-	-	-	-	1	-	1	-	-
UNIT	-	-	-	-	-	-	-	1	-	1	-	-
MONEY	1	1	-	-	-	-	-	-	-	-	-	1
CURR	1	1	-	-	-	-	-	-	-	-	-	1
<b>Overall</b>	<b>0.97</b>	<b>1</b>	<b>0.95</b>	<b>0.94</b>	<b>0.98</b>	<b>0.94</b>	<b>1</b>	<b>0.97</b>	<b>0.98</b>	<b>0.97</b>	<b>1</b>	<b>1</b>

Table 7: The F1-score Inter-Annotator Agreement between each pair of annotators (1-2 refers to F1-score between the first and second annotators)

#### 4.1. Multi-Task Learning for Nested NER

This section presents a multi-task based approach for nested NER (Figure 3). Similar to some of the methods we discussed in section 2, our approach assumes the sequence is tagged once with the same entity type.

The model consists of the sequence encoder and multiple classifiers, one for each entity type. We used transformer encoder based on the BERT language model. There are few BERT models pre-trained on Arabic text including AraBERT (Antoun et al., 2020) and ARBERT (Abdul-Mageed et al., 2021). AraBERT was trained on two major datasets, the 1.5 billion words Arabic Corpus (El-Khair, 2016) and the Open Source International Arabic News Corpus (OSIAN), which consists of 3.5 million articles ( $\sim 1$  billion tokens), from 31 news sources in 24 Arab countries (Zeroual et al., 2019). The final size of AraBERT dataset is 70M sentences, corresponding to about 24GB of text.

BERT takes a sequence of tokens as input and for each token outputs a 768 dimensional representation. BERT relies on the attention mechanism (Vaswani et al., 2017) to learn the contextual information of a given token from the entire sequence. The pre-trained BERT used in this paper is the AraBERT. Following the BERT layer are the classification layers, which are linear layers used to project the 768 dimensional token representation into the output space. The multi-task model consists of 21 classification layers, one layer for each entity type listed in Table 4. Since we use the IOB2 tagging scheme, each linear layer is a multi-class classifier that outputs the probability distribution through softmax activation function for three classes,  $C \in \{I, O, B\}$ . The model is trained with cross entropy loss objective computed for each linear layer separately, which are summed to compute the final cross entropy loss.

#### 4.2. Model Training

The dataset was divided into three different datasets: training (385K tokens, 70%), validation (55K tokens, 10%) and test (110K tokens, 20%). We fine-tuned AraBERT (Antoun et al., 2020) on the nested NER tasks with learning rate,  $\eta = 1e^{-3}$ , batch size of 32, maximum of 20 epochs with early termination if there is no improvement on the validation data after five epochs. In general the model converged around epoch nine. The source code and the nested NER pre-trained model are available on GitHub<sup>3</sup>.

#### 4.3. Experiments and Results

We evaluated our model on Wjood, the dataset presented in this paper. However, there is one limitation to the model, it cannot recognize nested entities of same type. For example, in Table 5, the model can tag (اتحاد العاملين بجامعة القاهرة / *The Employee Union of the Cairo University*) with ORG, but cannot tag the sub-entity mention (جامعة القاهرة / *Cairo University*) with ORG. This is a limitation we plan to address in future work.

Overall, the model achieved 0.884 micro F1-score, 0.8772 precision and 0.8909 recall (Table 8). The entity type on which the model was most confident about is GPE with an F1-score of 0.947, while the model was least confident about QUANTITY, UNIT, PERCENT, WEBSITE and TIME with 0.2, 0.25, 0.4426, 0.4936 and 0.5526 F1-score, respectively. The reason for low accuracy on these entity types is the lack of sufficient annotation. As shown in Table 3, QUANTITY, UNIT and PERCENT have only 57, 59 and 137 mentions, respectively.

<sup>3</sup><https://ontology.birzeit.edu/wjood> (PIN=reviewer) the source code will be publicly available pending review decision.



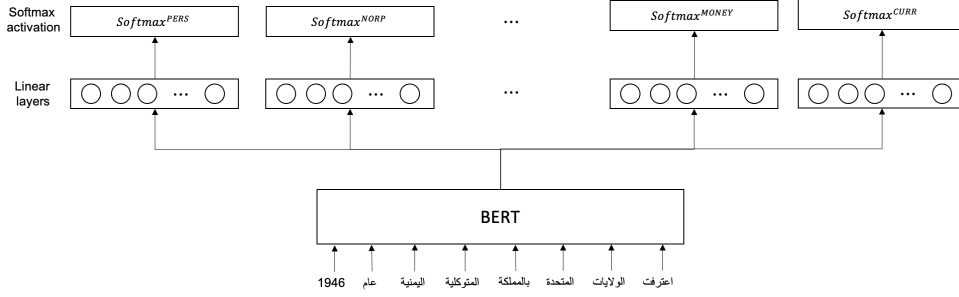


Figure 3: Nested NER model architecture

Tag	Precision	Recall	F1-Score
PERS	0.9135	0.9122	0.9129
NORP	0.6828	0.7037	0.6931
OCC	0.7993	0.8402	0.8193
ORG	0.8924	0.9072	0.8997
GPE	0.9424	0.9516	0.9470
LOC	0.8054	0.7059	0.7524
FAC	0.7366	0.6481	0.6895
PRODUCT	0.3333	0.2500	0.2857
EVENT	0.6364	0.6488	0.6425
DATE	0.9253	0.9394	0.9323
TIME	0.6000	0.5122	0.5526
LANGUAGE	0.9310	0.7105	0.8060
WEBSITE	0.4496	0.5472	0.4936
LAW	0.8525	0.9123	0.8814
CARDINAL	0.8437	0.8575	0.8505
ORDINAL	0.9411	0.9448	0.9430
PERCENT	0.2903	0.9310	0.4426
QUANTITY	0.2500	0.1667	0.2000
UNIT	0.5000	0.1667	0.2500
MONEY	0.9143	0.8205	0.8649
CURR	0.8810	0.9487	0.9136
<b>Overall</b>	<b>0.8772</b>	<b>0.8909</b>	<b>0.8840</b>

Table 8: Nested NER Results

## 5. Implementation

A RESTful web service for Arabic NER is developed and deployed online<sup>4</sup> as part of our language understanding resources (Jarrar and Amayreh, 2019; Al-Hajj and Jarrar, 2021b; Jarrar et al., 2019; Jarrar et al., 2021). The web service takes a text as input and returns the output in three different formats: (i) JSON IOB2, a JSON in which each token in the input text is returned with its corresponding tag similar to the IOB2 scheme, (ii) JSON entities, only the recognized named entities and their positions are returned, and (iii) XML, which is similar to the format (ii), but the named entities are marked up using XML. Additionally, as Figure 4 illustrates, a user interface is developed on top the web service for demonstration purposes, in which nested named entities are highlighted.

جامعة بيرزيت وبالتعاون مع مؤسسة ادوارد سعيد تنظم مهرجان للفن الشعبي، سيبدأ المهرجان الساعة الرابعة عصرا، بتاريخ 16/5/2022، وذلك برعاية من بنك فلسطين بمبلغ خمسة الاف دولار.

output format : highlighted

EVENT: تنظيم مهرجان للفن الشعبي, ORG: مؤسسة ادوارد سعيد, ORG: وبالتعاون مع, ORG: جامعة بيرزيت, GPE: سيبدأ المهرجان الساعة الرابعة عصرا, TIME: بتاريخ 16/5/2022, DATE: وذلك برعاية من بنك فلسطين, GPE: فلسطين, MONEY: بمبلغ, CURRENCY: دولار, ORG: خمسة الاف

Figure 4: Web user interface of the nested named entities.

## 6. Conclusions and Future Work

In this paper, we have presented Wojood, a new large corpus for Arabic nested NER, which consists of 550K tokens collected from MSA and dialect text in multiple domains, which was manually annotated using a rich set of 21 entity types. The IAA evaluation of the corpus, using  $Kappa$  and F1-score, illustrated a very high agreements between the annotators. We also trained a nested NER model based on multi-task learning using AraBERT, in which we achieved an F1-score of 0.884. We plan to enrich the corpus in several directions. We are working to increase the size of the corpus by including more specialized domains and new entity types. We also plan to manually link the extracted entities with nodes in the Wikidata Knowledge Graph in order to facilitate Arabic Entity Linking research. Another natural extension to this work is relationship extraction, but this annotating the relationships first. We will also address the limitation of our model architecture so that it can predict nested entities of the same type. Last but not least, we plan on linking the entities to concepts in the Arabic ontology (Jarrar, 2021; Jarrar, 2011) to enable richer semantic understanding of text.

## Acknowledgements

This research is partially funded by the Palestinian Higher Council for Innovation and Excellence. We would like to thank Taymaa Hammouda and Wasim Khatib for the technical support, and Prof. Fadi Zaraket for suggestions in the early phases. The authors also acknowledge the great efforts of many students who helped in the annotation process, especially Rania Shahwan, Donia Baninemra, Shimaa Hamayel, Ayah Hamdan, Jenan Morrar, Sondos Ilaiwi, and Alaa Omar.

<sup>4</sup> <https://ontology.birzeit.edu/wojood>



## 7. Bibliographical References

- Abdallah, S., Shaalan, K., and Shoaib, M. (2012). Integrating rule-based system with classification for arabic named entity recognition. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 311–322. Springer.
- Abdul-hamid, A. and Darwish, K. (2010). Simplified feature set for arabic named entity recognition. *Proceedings of the 2010 Named Entities Workshop*.
- Abdul-Mageed, M., Elmadany, A. R., and Nagoudi, E. M. B. (2021). Arbert & marbert: Deep bidirectional transformers for arabic. In *Proceedings of the ACL-IJCNLP 2021*. ACL.
- Abuaiadah, D., Rajendran, D., and Jarrar, M. (2017). Clustering arabic tweets for sentiment analysis. In *Proceedings of the 2017 IEEE/ACS 14th International Conference on Computer Systems and Applications*, pages 499–506. IEEE Computer Society, November.
- Al-Hajj, M. and Jarrar, M. (2021a). 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)*, pages 40–48, Online, sep. INCOMA Ltd.
- Al-Hajj, M. and Jarrar, M. (2021b). 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, aug. Association for Computational Linguistics.
- Al-Qurishi, M. and Souissi, R. (2021). Arabic named entity recognition using transformer-based-crf model. In *Proceedings of the ICNLSP 2021 Conference*.
- Ali, B. A. B., Mihi, S., Bazi, I. E., and Laachfoubi, N. (2020). A recent survey of arabic named entity recognition on social media. *Revue d'Intelligence Artificielle*, 34.
- Antoun, W., Baly, F., and Hajj, H. (2020). Arabert: Transformer-based model for arabic language understanding. In *LREC 2020 Workshop Language Resources and Evaluation Conference*, 2.
- Benajiba, Y., Rosso, P., and Ruiz, J. M. B. (2007). Anersys: An arabic named entity recognition system based on maximum entropy. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 4394 LNCS.
- Bhatia, P., Celikkaya, B., and Khalilia, M. (2019). Joint entity extraction and assertion detection for clinical text. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 954–959, Florence, Italy, July. Association for Computational Linguistics.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 06.
- Brandsen, A., Verberne, S., Wansleeben, M., and Lambers, K. (2020). Creating a dataset for named entity recognition in the archaeology domain. In *LREC 2020 - 12th International Conference on Language Resources and Evaluation, Conference Proceedings*.
- Campillos-Llanos, L., Valverde-Mateos, A., Capllonch-Carrión, A., and Moreno-Sandoval, A. (2021). A clinical trials corpus annotated with umls entities to enhance the access to evidence-based medicine. *BMC Medical Informatics and Decision Making*, 21.
- Dai, X. and Adel, H. (2020). An analysis of simple data augmentation for named entity recognition. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3861–3867, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Darwish, K., Habash, N., Abbas, M., Al-Khalifa, H., Al-Natsheh, H. T., Bouamor, H., Bouzoubaa, K., Cavalli-Sforza, V., El-Beltagy, S. R., El-Hajj, W., Jarrar, M., and Mubarak, H. (2021). A panoramic survey of natural language processing in the arab worlds. *Commun. ACM*, 64(4):72–81, April.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.
- Di Eugenio, B. and Glass, M. (2004). The kappa statistic: A second look. *Computational Linguistics*, 30(1):95–101.
- El-Khair, I. A. (2016). 1.5 billion words arabic corpus. *arXiv preprint arXiv:1611.04033*.
- Feng, X., Qin, B., and Liu, T. (2018). A language-independent neural network for event detection. *Science China Information Sciences*, 61.
- Finkel, J. R. and Manning, C. D. (2009). Nested named entity recognition. In *Proceedings of ACL and AFNLP 2009*, pages 141–150. ACL.
- Guo, J., Xu, G., Cheng, X., and Li, H. (2009). Named entity recognition in query. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 267–274.
- Haff, K. E., Jarrar, M., Hammouda, T., and Zaraket, F. (2022). Curras + baladi: Towards a levantine corpus. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2022)*, Marseille, France, June.
- Hassan, H. and Sorensen, J. (2005). An integrated approach for arabic-english named entity translation. In *Proceedings of the ACL Workshop on Computa-*

- tional Approaches to Semitic Languages*, pages 87–93.
- Hripcsak, G. and Rothschild, A. S. (2005). Agreement, the f-measure, and reliability in information retrieval. *Journal of the American Medical Informatics Association*, 12.
- Jaber, A. and Zaraket, F. A. (2018). Morphology-based entity and relational entity extraction framework for arabic. *Revue Traitement Automatique des Langues*, 58.
- Jarrar, M. and Amayreh, H. (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, June.
- Jarrar, M., Deik, A., and Faraj, B. (2011). Ontology-based data and process governance framework -the case of e-government interoperability in palestine. In *Proceedings of the IFIP International Symposium on Data-Driven Process Discovery and Analysis (SIMPDA'11)*, pages 83–98, 6.
- Jarrar, M., Habash, N., Akra, D., and Zalmout, N. (2014). Building a corpus for palestinian arabic: a preliminary study. In *Proceedings of the EMNLP 2014, Workshop on Arabic Natural Language*, pages 18–27. Association For Computational Linguistics, October.
- Jarrar, M., Habash, N., Alrimawi, F., Akra, D., and Zalmout, N. (2017). Curras: An annotated corpus for the palestinian arabic dialect. *Journal Language Resources and Evaluation*, 51(3):745–775, September.
- Jarrar, M., Amayreh, H., and McCrae, J. P. (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, May.
- Jarrar, M., Karajah, E., Khalifa, M., and Shaalan, K. (2021). Extracting synonyms from bilingual dictionaries. In *Proceedings of the 11th International Global Wordnet Conference (GWC2021)*, pages 215–222. Global Wordnet Association, Jan.
- Jarrar, M. (2011). Building a formal arabic ontology (invited paper). In *Proceedings of the Experts Meeting on Arabic Ontologies and Semantic Networks*. ALECSO, Arab League, 7.
- Jarrar, M. (2021). The arabic ontology - an arabic wordnet with ontologically clean content. *Applied Ontology Journal*, 16(1):1–26.
- Ju, M., Miwa, M., and Ananiadou, S. (2018). A neural layered model for nested named entity recognition. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1446–1459, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Katiyar, A. and Cardie, C. (2018). Nested named entity recognition revisited. In *Proceedings of the NAACL HLT 2018*, volume 1, pages 861–871. ACL.
- Khalifa, M. and Shaalan, K. (2019). Character convolutions for arabic named entity recognition with long short-term memory networks. *Computer Speech and Language*, 58.
- Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., and Dyer, C. (2016). Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California, June. Association for Computational Linguistics.
- Li, L., Jin, L., Jiang, Z., Song, D., and Huang, D. (2015). Biomedical named entity recognition based on extended recurrent neural networks. In *2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 649–652.
- Lu, W. and Roth, D. (2015). Joint mention extraction and classification with mention hypergraphs. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 857–867, Lisbon, Portugal, September. Association for Computational Linguistics.
- Martínez-deMiguel, C., Segura-Bedmar, I., Chacón-Solano, E., and Guerrero-Aspizua, S. (2022). The raredis corpus: A corpus annotated with rare diseases, their signs and symptoms. *Journal of Biomedical Informatics*, 125:103961, 1.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. In Yoshua Bengio et al., editors, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Mohit, B., Schneider, N., Bhowmick, R., Ofizer, K., and Smith, N. (2012). Recall-oriented learning of named entities in arabic wikipedia. In *EACL 2012 Proceedings*.
- Mostefa, D., Laïb, M., Chaudiron, S., Choukri, K., and Chalendar, G. D. (2009). A multilingual named entity corpus for arabic , english and french named entity tagset the afp corpus. In *Proceedings of the Second International Conference on Arabic Language Resources and Tools*.
- Ohta, T., Tateisi, Y., and Kim, J. (2002). The genia corpus: An annotated research abstract corpus in molecular biology domain. In *Proceedings of the second international conference on Human Language Technology Research*, pages 82–86. Morgan Kaufmann Publishers Inc.
- Passos, A., Kumar, V., and McCallum, A. (2014). Lexicon infused phrase embeddings for named entity resolution. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 78–86, Ann Arbor, Michigan, June. Association for Computational Linguistics.

- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Ringland, N., Dai, X., Hachey, B., Karimi, S., Paris, C., and Curran, J. R. (2020). Nne: A dataset for nested named entity recognition in english newswire. In *Proceedings of ACL 2019*. ACL.
- Salah, R. E. and Zakaria, L. Q. B. (2018). Building the classical arabic named entity recognition corpus (caner corpus). *Journal of Theoretical and Applied Information Technology*, 96.
- Sang, E. F. T. K. and Veenstra, J. (1999). Representing text chunks. In *Proceedings of the Ninth Conference on European Chapter of the Association for Computational Linguistics*, EACL '99, page 173–179, USA. Association for Computational Linguistics.
- Shaalán, K. and Raza, H. (2007). Person name entity recognition for arabic. In *Proceedings of the 2007 Workshop on Computational Approaches to Semitic Languages: Common Issues and Resources*, pages 17–24.
- Shaalán, K. (2014). A survey of arabic named entity recognition and classification. *Computational Linguistics*, 40.
- Shaheen, M. and Ezzeldin, A. M. (2014). Arabic question answering: systems, resources, tools, and future trends. *Arabian Journal for Science and Engineering*, 39(6):4541–4564.
- Shibuya, T. and Hovy, E. (2020). Nested Named Entity Recognition via Second-best Sequence Learning and Decoding. *Transactions of the Association for Computational Linguistics*, 8:605–620, 09.
- Sohrab, M. G. and Miwa, M. (2018). Deep exhaustive model for nested named entity recognition. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2843–2849, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Tjong Kim Sang, E. F. and De Meulder, F. (2003). Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Tjong Kim Sang, E. F. (2002). Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition. In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Uzuner, Ö., South, B. R., Shen, S., and DuVall, S. L. (2011). 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, ., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Walker, C., Strassel, S., Medero, J., and Maeda, K. (2006). Ace 2005 multilingual training corpus ldc2006t06. Technical report, Linguistic Data Consortium, Philadelphia.
- Wang, J., Shou, L., Chen, K., and Chen, G. (2020). Pyramid: A layered model for nested named entity recognition. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5918–5928, Online, July. Association for Computational Linguistics.
- Weischedel, R., Pradhan, S., Ramshaw, L., Palmer, M., Xue, N., Marcus, M., Taylor, A., Greenberg, C., Hovy, E., Belvin, R., et al. (2011). Ontonotes release 4.0. *LDC2011T03*, Philadelphia, Penn.: Linguistic Data Consortium.
- Weischedel, R., Palmer, M., Marcus, M., Hovy, E., Pradhan, S., Ramshaw, L., Xue, N., Taylor, A., Kaufman, J., Franchini, M., et al. (2013). Ontonotes release 5.0 ldc2013t19. Technical report, Linguistic Data Consortium.
- Zeroual, I., Goldhahn, D., Eckart, T., and Lakhouaja, A. (2019). OSIAN: Open source international Arabic news corpus - preparation and integration into the CLARIN-infrastructure. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 175–182, Florence, Italy, August. Association for Computational Linguistics.