

TURJUMAN: A Public Toolkit for Neural Arabic Machine Translation

El Moatez Billah Nagoudi* AbdelRahim Elmadany* Muhammad Abdul-Mageed*

Deep Learning & Natural Language Processing Group

The University of British Columbia

{moatez.nagoudi,a.elmadany,muhammad.mageed}@ubc.ca

Abstract

We present TURJUMAN, a neural toolkit for translating from 20 languages into Modern Standard Arabic (MSA). TURJUMAN exploits the recently-introduced text-to-text Transformer AraT5 model, endowing it with a powerful ability to decode into Arabic. The toolkit offers the possibility of employing a number of diverse decoding methods, making it suited for acquiring paraphrases for the MSA translations as an added value. To train TURJUMAN, we sample from publicly available parallel data employing a simple semantic similarity method to ensure data quality. This allows us to prepare and release AraOPUS-20, a new machine translation benchmark. We publicly release our translation toolkit (TURJUMAN) as well as our benchmark dataset (AraOPUS-20).¹

Keywords: Machine Translation, Neural Machine Translation, Arabic, Arabic NLP, Open Source, TURJUMAN, Toolkit.

1. Introduction

Natural language processing (NLP) technologies such as question answering, machine translation (MT), summarization, and text classification are witnessing a surge. This progress is the result of advances in deep learning methods, availability of large datasets, and increasingly powerful computing infrastructure. As these technologies continue to mature, their applications in everyday life become all the more pervasive. For example, neural machine translation (NMT), the focus of the current work, has applications in education, health, tourism, search, security, recreation, etc. Similar to other areas, progress in NMT is contingent on high-quality, standardized datasets and fast prototyping. Such datasets and tools are necessary for meaningful comparisons of research outcomes, benchmarking, and training of next generation scholars. Off-the-shelf tools are also especially valuable both as stand-alone and enabling technologies in all research and development. Although various tools have been developed for Arabic NLP tasks such as those involving morphosyntactic analysis (Pasha et al., 2014; Darwish and Mubarak, 2016; Obeid et al., 2020b) and detection of social meaning (Abdul-Mageed et al., 2019; Farha and Magdy, 2019), there has not been as much progress for MT. More specifically, there is shortage of publicly available tools for Arabic MT. The goal of this work is to introduce **TURJUMAN**, a new publicly available Arabic NMT toolkit that seeks to contribute to bridging this gap.

Recent advances in NMT leverages progress in Transformer-based encoder-decoder language models, and TURJUMAN takes advantage of such a progress. In particular, encoder-decoder models such as MASS (Song et al., 2019), BART (Lewis et al., 2020), and T5 (Raffel et al., 2019), and their multilingual counter-parts have all

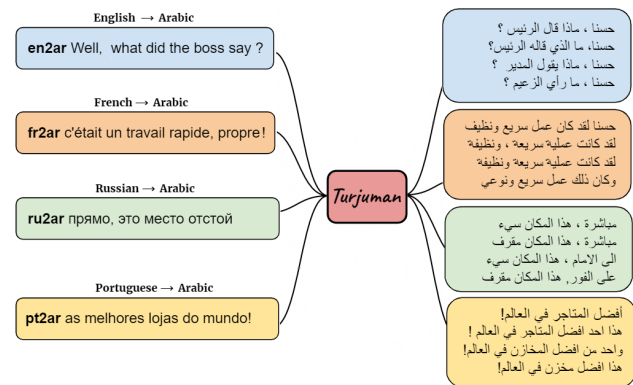


Figure 1: Our TURJUMAN neural machine translation toolkit illustrated with four prompt MT tasks: English, French, Russian, and Portuguese \rightarrow Arabic. For each source sentence, we employ four decoding methods to produce output: *greedy search*, *beam search*, *top-k sampling*, and *top-p sampling*.

been shown to remarkably benefit NMT. For this reason, TURJUMAN is built off AraT5 (Nagoudi et al., 2022). AraT5 is a recently released text-to-text Transformer model. For comparisons, we benchmark AraT5 against a number of baselines on a new parallel dataset that we also introduce as part of this work. Importantly, we do not intend TURJUMAN as a tool for delivering state-of-the-art (SOTA) translations. For this reason, we do not use all parallel datasets at our disposal. Rather, we introduce TURJUMAN as an extensible framework. For example, it can be further developed to produce SOTA performance by fine-tuning its backend model on larger datasets.

In the context of creating our tool, we also prepare and release **AraOPUS-20**. AraOPUS-20 is a reasonably-sized parallel dataset of 20 language pairs (with X \rightarrow Arabic) for NMT. We extract AraOPUS-20 from OPUS (Tiedemann, 2012). Since OPUS is known to involve noisy translations, we propose a simple quality assurance method based on semantic similarity to remove this noise from the dataset. We

¹<https://github.com/UBC-NLP/turjuman>

*All authors contributed equally.

release AraOPUS-20 in standard splits, thereby making it well-suited for Arabic MT model comparisons.

TURJUMAN also integrates recent progress in diverse decoding, such as *greedy search* (Cormen et al., 2009), *beam search* (Koehn, 2009), *top-k sampling* (Fan et al., 2018), and *nucleus sampling* (Holtzman et al., 2019). This makes it possible to use TURJUMAN for generating various translations of the same foreign sequence. As such, TURJUMAN can also be used for producing paraphrases at the Arabic side (see Figure 1).

To summarize, we make the following contributions:

1. We prepare and release *AraOPUS-20*, an MT benchmark that we extract from the freely available parallel corpora OPUS (Tiedemann, 2012). AraOPUS-20 consists of bitext between Arabic and 20 languages. The languages paired with Arabic include high-resource languages such as English, French, and Spanish and low-resource ones such as Cebuano,¹ Tamashek,² and Yoruba³.
2. We introduce *TURJUMAN*, a python-based NMT toolkit for translating sentences from 20 languages into Arabic. TURJUMAN fine-tunes AraT5 (Nagoudi et al., 2022), a powerful Arabic text-to-text Transformer language model. Our toolkit can be used off-the-shelf as a strong baseline, or as an enabling technology. It is also extensible. For example, it can be further developed through additional fine-tuning on larger amounts of data.
3. We endow TURJUMAN with a diverse set of decoding capabilities, making it valuable for generating paraphrases (Fadaee et al., 2017) of foreign content into Arabic.

The rest of the paper is organized as follows: We provide an overview of works related to Arabic machine translation in Section 2. We introduce AraOPUS-20 MT benchmark in Section 3. We describe TURJUMAN in Section 4., and Section 6. is where we conclude.

2. Related Work

Our work is related to research on MT datasets and tools, and language models on which these tools may be fine-tuned. Hence, we start our coverage of related work by presenting most of the popular Arabic MT datasets for both MSA and Arabic dialects. We then provide an overview of Arabic MT systems and tools. Finally, we review both Arabic and multilingual encoder-decoder pre-trained language models since these are most relevant to the translation task.

¹Language spoken in the southern Philippines

²Tamashek or Tamasheq is a variety of Tuareg, a Berber macro-language widely spoken by nomadic tribes across North Africa countries.

³Yoruba is a language spoken in West Africa, primarily in Southwestern Nigeria.

2.1. MSA MT Resources

Open Parallel Corpus (OPUS). Tiedemann (2012) propose the large, multi-lingual, parallel sentences datasets OPUS. OPUS contains more than 2.7 billion parallel sentences in 90 languages including Arabic.⁴ We extract AraOPUS-20 from OPUS. A number of additional MSA datasets involving Arabic have also been proposed. Although we do not make use of any of these, we review them here both for completeness and since they can be exploited for extending TURJUMAN.

United Nations Parallel Corpus. Ziemski et al. (2016) introduce a manually translated united nations (UN) documents corpus covering the six official UN languages: *Arabic, Chinese, English, French, Russian, and Spanish*. The corpus consists of development and test sets only, each of which comprise 4K sentences that are one-to-one alignments across all official languages.

IWSLT Corpus. Several Arabic to English parallel datasets were released during IWSLT⁵ evaluation campaigns. These include IWSLT 2012 (Federico et al., 2012), IWSLT 2013 (Cettolo et al., 2013), IWSLT 2016 (Cettolo et al., 2016), and IWSLT 2017 (Cettolo et al., 2017).

Arab-Acquis. Habash et al. (2017) propose Arab-Acquis. It consists of 12k English and French sentences extracted from the JRC-Acquis corpus (Steinberger et al., 2006). The foreign sentences are translated into Arabic by two professional translators. JRC-Acquis is a publicly available parallel collection of legislative text of the European Union and is written in the 22 official European languages.

MSA MADAR Corpus. Proposed by Bouamor et al. (2018), this dataset. It consists of 10k MSA-English evaluation sentences manually translated using the crowdsourcing platform *crowdFlower.com*.⁶

2.2. Dialectal MT Resources

There are also a number of available dialectal datasets that can be used to extend TURJUMAN. We also briefly review these here.

APT Corpus. Zbib et al. (2012) present an Arabic-English dataset⁷ covering MSA and two other Arabic dialects. It comprises 8.11M MSA-English sentences, 138k Levantine-English sentences, and 38k Egyptian-English sentences. The dataset was collected from Arabic weblogs and translation was carried out through Amazon Mechanical Turk.⁸

Qatari-English Speech Corpus. This parallel corpus comprises 14.7k Qatari-English sentences collected by Elmahdy et al. (2014) from talk-show programs and Qatari TV series and translated into English.

Multi-dialectal Parallel Corpus (MDPC). Bouamor et al. (2014) construct MDPC by selecting 2k Egyptian-English sentences from the APT corpus (Zbib et al., 2012). Then, native speakers from Palestine, Syria, Jordan, and Tunisia

⁴<https://opus.nlpl.eu/>

⁵<https://wit3.fbk.eu>.

⁶<http://www.crowdflower.com/>.

⁷<https://catalog.ldc.upenn.edu/LDC2012T09>.

⁸<https://www.mturk.com>.

MSA sentences	English sentences	Sim
ألمانيا	The annex to the present report.	-0.12
إساءة استغلال مركز الهيمنة	Abuse of dominance	0.02
البند ٣ (أ) من جدول الأعمال	Draft resolution submitted by Argentina	0.11
ألف - آراء بشأن تنظيم الدورة	plenary meetings have been made.	0.27
نفس الوطن .	the same home, yes! we are brothers!	0.43
خامساً اعتماد التقرير	Adoption of the report	0.52

ألبانيا والكاميرون.	Albania, Cameroon.	0.70
مشاورات غير رسمية	Informal consultation	0.72
نوارة نجم تقع في حب أحمد مكي:	Nawara Negm is falling in love with Mekki	0.74
نتيجة التصويت كما يلي:	The voting was as follows:	0.77
المجموعة العربية	Arab Group	0.91
وتشمل هذه التكاليف ، تكاليف الصيانة	These costs include, maintenance	0.94

Table 1: A sample of MSA-English parallel sentences extracted from the Open Parallel Corpus (OPUS). We report semantic similarity on each pair of sentences using the multilingual sentence transformer model SBERT. **Green:** Selected sentences. **Red:** Ignored sentences.

were asked to translate the sentences into their respective native dialects.

Parallel Arabic Dialect Corpus (PADIC). Meftouh et al. (2015) offers PADIC, a multi-dialect corpus including MSA, Algerian, Tunisian, Palestinian, and Syrian. PADIC consists of 6.4K parallel sentences between MSA and all the listed dialects.

Dial2MSA. Mubarak (2018) release this parallel dialectal Arabic corpus for converting dialectal Arabic to MSA. The dataset has 6K tweets from four Arabic dialects: Egyptian, Levantine , Gulf ,and Maghrebi. Each of the dialects is translated into MSA by native speakers of each dialect.

DIA MADAR Corpus. Bouamor et al. (2018) introduce this commissioned corpus. Arabic native speakers from 25 Arabic cities were tasked to translate 2k English sentences each into their own native dialect. The sentences are selected from the Basic Traveling Expression Corpus (Takezawa et al., 2007). We now review systems for Arabic MT.

2.3. Arabic MT Systems

MSA MT. Arabic MT went through different stages, including rule-based systems (Bakr et al., 2008; Mohamed et al., 2012; Salloum and Habash, 2013) and statistical MT (Habash and Hu, 2009; Salloum and Habash, 2011; Ghoneim and Diab, 2013). There has been work on Arabic MT employing neural methods. For example, Almahairi et al. (2016) propose an Arabic \leftrightarrow English NMT using a vanilla attention-based NMT model of Bahdanau et al. (2014). Also, Junczys-Dowmunt et al. (2016) report an experimental study where phrase-based NMT across 30 translation directions including Arabic is investigated. Other sentence-based Arabic \leftrightarrow English NMT systems

training on various datasets are presented in Akeel and Mishra (2014), Durrani et al. (2017), and Alrajeh (2018). A number of Arabic-related NMTs were also proposed to translate from languages other than English into Arabic. This includes from Chinese (Aqlan et al., 2019), Turkish (El-Kahlout et al., 2019), Japanese (Noll et al., 2019), and four foreign languages⁹ into MSA (Nagoudi et al., 2022).

Dialectal MT. Some work has focused on translating between MSA and Arabic dialects. For instance, Zbib et al. (2012) show the impact of combined MSA and dialectal data on dialect/MSA \rightarrow English MT performance. Sajjad et al. (2013) use MSA as a pivot language for translating Arabic dialects into English. Salloum et al. (2014) investigate the effect of sentence-level dialect identification and several linguistic features for dialect/MSA \rightarrow English translation. Guellil et al. (2017) propose an NMT system for Arabic dialects using a vanilla recurrent neural network encoder-decoder model for translating Algerian Arabic written in a mixture of Arabizi and Arabic characters into MSA. Baniata et al. (2018) present an NMT system to translate Levantine (Jordanian, Syrian, and Palestinian) and Maghrebi (Algerian, Moroccan, and Tunisia) into MSA. Sajjad et al. (2020) introduce AraBench, an evaluation benchmark for dialectal Arabic to English MT and several NMT systems using several training settings: fine-tuning, data augmentation, and back-translation. Farhan et al. (2020) propose an unsupervised dialectal NMT where the source dialect is not represented in training data (i.e., zero-shot MT (Lample et al., 2018)). More recently, Nagoudi et al. (2021) introduce a transformer-based MT system for translating from code-

³ ⁹English, French, German, and Russian.

mixed Modern Standard Arabic and Egyptian Arabic into English. Nagoudi et al. (2022) Finally, propose three Arabic text-to-text transformer (AraT5) models dedicated to MSA and a diverse set of Arabic dialects. The models are used in several dialects \rightarrow English MT tasks. To the best of our knowledge, neither MSA nor dialectal machine translation systems described in this section have been made publicly available for research.

2.4. Open Source Arabic Tools

There have been many efforts to develop tools to support Arabic NLP. Some tools target morphosyntax such as in morphological analysis, disambiguation, POS tagging, and diacritization (Pasha et al., 2014; Darwish and Mubarak, 2016; Obeid et al., 2020a), while others focus on social meaning tasks such as sentiment analysis, emotion, age, gender, and sarcasm detection (Farha and Magdy, 2019; Abdul-Mageed et al., 2019). For MT, we do not know of any publicly available Arabic MT tools (let alone ones that afford many-to-Arabic translations nor diverse decoding). We now review Transformer-based Arabic and multilingual encoder-decoder models since these can be fine-tuned for MT.

2.5. Pre-Trained Language Models

mBART50 (Liu et al., 2020) is a multilingual encoder-decoder model primarily intended for MT. It is pre-trained by denoising full texts in 50 languages, including Arabic. Then, mBART is fine-tuned on parallel MT data under three settings: many-to-English, English-to-many, and many-to-many. The parallel MT data used contains a total of 230M parallel sentences and covers high-, mid-, and low-resource languages.

mT5 (Xue et al., 2020) is the multilingual version of Text-to-Text Transfer Transformer model (T5) (Raffel et al., 2019). The basic idea behind this model is to treat every text-based language task as a “text-to-text” problem, (i.e. taking text format as input and producing new text format as output), where a multi-task learning set-up is applied to several NLP tasks: question answering, document summarization, and MT. The mT5 model is pre-trained on the “mC4: Multilingual Colossal Clean Crawled Corpus”, which is \sim 26.76TB for 101 languages (including Arabic). **AraT5** (Nagoudi et al., 2022) is an Arabic text-to-text Transformer model dedicated to MSA and Arabic dialects. Again, AraT5 is an encoder-decoder Transformer similar in configuration and size to T5 (Raffel et al., 2019). AraT5 was trained on more than 248GB of Arabic text (70GB MSA and 178GB tweets). In addition to Arabic, AraT5’s vocabulary covers 11 others languages. Namely, the model covers vocabulary from Bulgarian, Czech, English, French, German, Greek, Italian, Portuguese, Russian, Spanish, and Turkish.

3. AraOPUS-20 Parallel Dataset

In this section, we describe AraOPUS-20 (the dataset we use to develop TURJUMAN) and the cleaning process we employ to ensure high quality of the data.

xx \rightarrow ar	Orig. OPUS	Filtering	Train	Dev	Test
bg	2M	sim	1M	2K	2K
bs	2M	rand	1M	2K	2K
cs	2M	sim	1M	2K	2K
da	2M	sim	0.93M	2K	2K
de	2M	sim	0.99M	2K	2K
el	2M	sim	1M	2K	2K
en	2M	sim	1M	2K	2K
es	2M	sim	1M	2K	2K
fr	2M	sim	1M	2K	2K
hi	2M	sim	0.81M	2K	2K
it	2M	sim	1M	2K	2K
ko	2M	sim	0.83M	2K	2K
pl	2M	rand	1M	2K	2K
pt	2M	sim	1M	2K	2K
ru	2M	sim	1M	2K	2K
tr	2M	sim	1M	2K	2K
ceb	83.1K	all	82.1K	200	200
gd	19.9K	all	19.5K	200	200
tmh	2.6K	all	2.8K	100	100
yo	1.4K	all	1.2K	100	100

Table 2: OPUS filtering process and data distribution in AraOPUS-20. Filtering methods: **(1) sim**: keep only 1M sentences with semantic similarity between [0.7,0.99] **(2) random**: if the language is not supported by SBERT model, we pick a random 1M pair of sentences. **(3) all**: for the low resource languages, we keep all the parallel data. **ar**: Arabic **bg**: Bulgarian. **bs**: Bosnian. **cs**: Czech. **da**: Danish. **de**: German. **el**: Greek. **en**: English. **es**: Spanish. **fr**: French. **hi**: Hindi. **it**: Italian. **ko**: Korean. **pl**: Polish. **pt**: Portuguese. **ru**: Russian. **tu**: Turkish. **ceb**: Cebuano. **gd**: Scots Gaelic. **tmh**: Tamashek. **yo**: Yoruba.

3.1. Training Data.

As mentioned earlier, in order to develop our TURJUMAN tool, we use AraOPUS-20. AraOPUS-20 is extracted from OPUS (Tiedemann, 2012) as follows:

1. We randomly pick 2M Arabic parallel sentences from the 16 highest-resource languages from among our 20 languages. Namely, we extract parallel data involving Arabic (mainly MSA) and Bulgarian, Czech, English, French, German, Greek, Italian, Portuguese, Russian, Spanish, Hindi, Polish, Korean, and Turkish.
2. We also use available data form the four low resource languages: Cebuano (Philippine), Scots Gaelic (Scotland), Tamashek (Mali), and Yoruba (Nigeria).¹⁰

3.2. Quality of Parallel Data.

In order to investigate the quality of OPUS Arabic parallel sentences, we measure semantic similarity between the parallel sentences by running a multilingual sentence Transformer model (Reimers and Gurevych, 2020) on each pair of sentences,¹¹ keeping only pairs with a semantic

¹⁰We list the country where a language is mostly spoken, otherwise a given language can be spoken in more than one country.

¹¹We exclude the low-resource data from our semantic similarity steps as these languages are not supported in Reimers and Gurevych (2020) model.

similarity score between 0.70 and 0.99. This allows us to filter out sentence pairs whose source and target are identical (i.e., similarity score = 1) and those that are not good translations of one another (i.e., those with a cross-lingual semantic similarity score < 0.70). Manually inspecting the data, we find that a threshold of $> 0.70\%$ safely guarantees acquiring semantically similar (i.e., good translations) and distinct pairs of sentences (i.e., sentences from two different languages). Table 1 shows a sample of MSA-English parallel sentences extracted from OPUS, along with their measured semantic similarity. We pick the top 1M sentences¹² (i.e., sentences with high semantic similarity score and satisfy our semantic similarity condition) from each language. We then split the resulting dataset into Train, Dev, and Test (see next section) and refer to the resulting benchmark that covers 20 languages as *AraOPUS-20* as we explained.

3.3. Development and Test Data.

For each of development and test split, we randomly pick 2k sentences from AraOPUS-20 (after filtering). We do this for all of the high resource languages. Regarding the low-resources languages, if the training split has more than 15K sentences, we randomly pick 200 sentences each for Dev and Test. Otherwise, we consider only 100 sentences for each of these splits per language. More details about the AraOPUS-20 parallel data distribution are given in Table 2.

4. TURJUMAN Tool

TURJUMAN is a publicly available toolkit for translating sentences from 20 languages into MSA. The package consists of a Python library and related command-line scripts. In this section, we discuss: (1) the training and evaluation processes of TURJUMAN’s backbone MT model and (2) how we design the TURJUMAN tool itself and its different settings.

4.1. Approach

Training. For all the 20 languages, we fine-tune AraT5 (Nagoudi et al., 2022) with training data of AraOPUS-20 (see in Section 3.). That is, we train a single *multilingual* model that translates from a given foreign language into MSA (many-to-MSA). Currently, a user needs to specify the identity of the input language.¹³ We train our models on 96 AMD MI50 GPUs (16GB each) for 25 epochs with a batch size of 32, maximum sequence length of 256 tokens, and a learning rate of $5e^{-5}$.

Evaluation. In order to evaluate our TURJUMAN model, we use two datasets (AraOPUS-20 and United Nations Parallel Corpus, both described in Section 3.3.).¹⁴ As a

¹²For low resource languages we use all available sentences.

¹³We plan to incorporate a language ID module into TURJUMAN in the future.

¹⁴We exclude the Chinese language as it is not included in our training data.

rule, for all datasets we identify the best model on our Dev data¹⁵ and blind-test it on our Test split for each language separately. For the two datasets, we report results on both Dev and Test splits as shown in Tables 5 and 6 respectively.

Baselines. For comparison, we use three baselines:

- **Baseline I.** A vanilla sequence-to-sequence (S2S) Transformer (Vaswani et al., 2017) as implemented in Fairseq (Ott et al., 2019). We train this model from scratch using AraOPUS-20 training data
- **Baseline II.** We fine-tune the multilingual encoder-decoder model mT5 (Xue et al., 2020) on the same training data as our second baseline.
- **Baseline III.** We use the mBART-50 many-to-many multilingual MT model for our third baseline. We do not fine-tune this model on AraOPUS-20 Train data as it is a checkpoint of mBART-large-50 (Liu et al., 2020) already fine-tuned on a multilingual MT dataset covering high-, mid-, and low-resource languages. In total, this model is fine-tuned with 230M parallel sentences from these 50 languages.

4.2. Implementation

We distribute TURJUMAN as a modular toolkit built around standard libraries including PyTorch (Paszke et al., 2019) and HuggingFace (Lhoest et al., 2021).

Command-Line Tools. We provide several command-line tools for translation and evaluation:

- *turjuman_interactive*: This interactive command line facility can be used for quick sentence-by-sentence translation exploiting our fine-tuned NMT model (TURJUMAN’s backbone translation model).
- *turjuman_translate*: This is the same as the interactive command. However, a path to a file containing source sentences is required.
- *turjuman_score*: This evaluates an output translation (output translations) against reference translation(s) in terms of a BLEU score.

4.3. Decoding Support

We also endow TURJUMAN with support for MT-based *paraphrase* generation by adding four decoding methods at the decoder side. We implement a number of prominent decoding methods used in the literature. Namely, we implement *greedy search*, *beam search* (Koehn, 2009), *top-k sampling* (Fan et al., 2018), and *nucleus sampling* (Holtzman et al., 2019). Table 3 shows example translations with TURJUMAN exploiting each of the four decoding methods. We now briefly describe each of these methods.

5 ¹⁵We merge all the development data for all the 20 languages.

# Decoding	Source/Translated Sentences
en2ar:	<i>She sort of grew up in front of everyone in Arkansas. Then as the spokesman for President Trump</i>
GS	لقد نشأت في وسط كل شخص في أركنساس ، ثم أصبحت متحدثة باسم الرئيس ترامب.
BS	لقد نشأت أمام الجميع في أركنساس ، ثم أصبحت متحدثة باسم الرئيس ترامب.
Top-k	لقد نشأت في وسط كل شخص في أركنساس ، ثم أصبحت متحدثة باسم الرئيس ترامب.
Top-p	ترعرعت ، بجانب كل شخص في أركنساس ، ثم أصبحت الناطق باسم الرئيس ترمب ،
fr2ar:	<i>Match Algérie et Cameroun : la grande victoire des fake news.</i>
GS	مباراة الجزائر والكاميرون : النصر الكبير للأخبار الكاذبة.
BS	مباراة الجزائر والكاميرون : الانتصار الكبير للأخبار المزيفة.
Top-k	مباراة الجزائر والكاميرون : النصر الكبير للأخبار الكاذبة.
Top-p	مباراة الجزائر والكاميرون : البطولة الرئيسة للأخبار المزيفة.
pt2ar:	<i>Já o governo federal não explicou os motivos da manutenção dos contratos</i>
GS	لم تعد الحكومة الاتحادية تفسر أسباب استمرار العقود.
BS	لم تعد الحكومة الاتحادية تفسر أسباب الإبقاء على العقود.
Top-k	لم تعد الحكومة الاتحادية تفسر أسباب استمرار العقود.
Top-p	لم تعد الحكومة الاتحادية تشرح أسبابه المتعلقة بالمحافظة على العقود بعد الآن.
ru2ar :	<i>Резкий скачок цен на зерновые из-за войны: что будет с ценами на продовольствие?</i>
GS	ارتفاع أسعار الحبوب بسبب الحرب : ماذا عن أسعار الغذاء ؟
BS	ارتفاع حاد في أسعار الحبوب بسبب الحرب : ماذا عن أسعار الغذاء ؟
Top-k	ارتفاع أسعار الحبوب بسبب الحرب : ماذا عن أسعار الغذاء ؟
Top-p	تناقص حاد في أسعار الأومونيا بسبب الحرب : ماذا لو تناقصت أسعار الغذاء ؟
tr2ar :	<i>Türkiye ile Ermenistan arasında, son dönemde yeniden başlayan doğrudan uçuşlardan birindeyiz.</i>
GS	نحن في رحلة طيران مباشرة بين تركيا وأرمينيا ، التي بدأت في الآونة الأخيرة.
BS	بين تركيا وأرمينيا ، نحن في واحدة من الرحلات الجوية المباشرة التي بدأت من جديد خلال الفترة الأخيرة.
Top-k	نحن في رحلة طيران مباشرة بين تركيا وأرمينيا ، التي بدأت في الآونة الأخيرة
Top-p	حيث أنه يتم نقل معظم الركاب الطيارين ضمن عصر النقل الجوي الداخلي.

Table 3: A sample of sentences from five foreign languages along with their MSA translations using four decoding methods. **GS:** Greedy Search. **BS:** Beam Search. **Top-k:** top-k sampling. **Top-p:** top-p sampling. **en2ar:** English to Arabic. **fr2ar** French to Arabic. **pt2ar:** Portuguese to Arabic. **ru2ar:** Russian to Arabic. **tr2ar:** Turkish to Arabic.

Greedy Search. Is a simple heuristic strategy aiming to select the word with the highest conditional probability as its next word at each timestep as shown in Formula (1):

$$w_t = \operatorname{argmax}_w P(w|w_{1:t-1}) \quad (1)$$

Beam Search. Beam search is an improved version of greedy search that uses a hyper-parameter *num_beams*. It is based on exploring the solution space and reduces the risk of missing hidden high probability word sequences by keeping the most likely *num_beams* of hypotheses sequences.

Top-k Sampling. A probabilistic decoding method proposed by Fan et al. (2018) that aims to avoid repetitions during decoding. This method also increase diversity of the output by using a simple, yet powerful sampling stochastic scheme called **top-k** sampling. First, the top *k* words with

the highest probability are selected. Then, we sample from this shortlist of words. This allows the other high-scoring tokens a chance of being picked. Formula (2) describe top-k sampling, where $V^{(k)}$ is the top-k vocabulary.

$$w_t = \sum_{w \in V^{(k)}} P(w|w_{1:t-1}) \quad (2)$$

Top-p Sampling. Also called *nucleus-sampling*, this method is proposed by Holtzman et al. (2019). It shares the same principle as the *top-k* method, and the only difference between the two is that Top-p sampling chooses from the smallest possible set of words the sum of whose probability is greater than a certain probability *p* (i.e., threshold). This method is described in Formula (3), where $V^{(p)}$ is the top-p vocabulary.

$$w_t = \sum_{w \in V^{(p)}} P(w|w_{1:t-1}) \geq p \quad (3)$$

Arguments	Description	Command-line	Required
<code>--help</code> or <code>-h</code>	show the help message and exit	interactive, translate	no
<code>--text</code> or <code>-t</code>	translate the input text into Arabic	translate	yes
<code>--input_file</code>	path of input file	translate	yes
<code>--batch_size</code> or <code>-bs</code>	the number of sentences translated in one iteration	translate	no
<code>--seq_length</code> or <code>-s</code>	generate sequences of maximum length <code>seq_length</code>	interactive, translate	no
<code>--search_method</code> or <code>-m</code>	decoding method [‘greedy’, ‘beam’, ‘sampling’]	interactive, translate	no
<code>--n_beam</code>	beam search with a size of <code>n_beam</code>	interactive, translate	no
<code>--top_k</code> or <code>-k</code>	sampling using <code>top-k</code>	interactive, translate	no
<code>--top_p</code> or <code>-p</code>	sampling using <code>top-p</code>	interactive, translate	no
<code>--no_repeat_ngram_size</code>	ngram size cannot be repeated in the generation	interactive, translate	no
<code>--max_outputs</code> or <code>-o</code>	number of hypotheses to output	interactive, translate	no
<code>--cache_dir</code> or <code>-c</code>	path of the cache directory	interactive, translate	no
<code>--logging_file</code> or <code>-l</code>	the logging file path	interactive, translate	no
<code>--hyp_file</code> or <code>-p</code>	path of hypothesis file	score	yes
<code>--ref_file</code> or <code>-g</code>	path of references file	score	yes

Table 4: *Required* and *optional* arguments for each of TURJUMAN command-line tools.

Model	Split	bg	bs	cs	da	de	el	en	es	fr	hi	it	ko	pl	pt	ru	tr	ceb	gd	tmh	yo
S2SMT	Dev	5.45	4.14	4.73	5.04	4.28	6.17	7.08	6.42	7.64	4.45	5.27	3.85	6.59	7.24	5.41	4.27	1.67	0.13	0.23	2.59
	Test	5.25	4.44	4.38	5.94	4.61	6.77	7.42	6.22	6.98	4.59	5.24	3.58	6.93	7.15	5.65	4.25	1.98	0.17	0.025	2.48
mBART	Dev	-	-	0.98	-	0.71	-	8.02	1.23	1.45	0.39	0.22	0.47	1.18	1.82	1.93	2.03	-	-	0.02	-
	Test	-	-	1.03	-	0.88	-	8.38	1.36	1.66	0.40	0.32	0.39	1.38	1.86	1.78	2.15	-	-	0.06	-
mT5	Dev	8.13	6.63	9.71	10.94	16.64	<u>14.28</u>	25.41	21.12	19.04	<u>4.29</u>	15.17	<u>6.08</u>	<u>3.29</u>	10.96	26.63	10.84	10.23	<u>2.37</u>	<u>0.58</u>	3.45
	Test	12.85	6.60	7.79	10.94	14.90	15.33	25.24	21.12	20.74	4.80	12.90	8.27	6.41	21.13	27.94	11.77	7.53	3.20	0.43	5.24
TURJ	Dev	<u>8.68</u>	<u>7.94</u>	<u>9.83</u>	<u>11.30</u>	<u>16.84</u>	13.82	<u>25.80</u>	<u>21.57</u>	<u>21.43</u>	2.86	<u>16.99</u>	2.18	3.43	<u>12.00</u>	<u>29.67</u>	<u>11.75</u>	<u>9.89</u>	2.32	0.64	<u>5.11</u>
	Test	13.64	7.87	8.32	11.30	16.05	15.06	25.46	21.57	22.43	3.29	14.92	3.44	6.38	23.64	31.68	13.05	8.43	2.41	0.29	4.39

Table 5: Results of TURJUMAN in BLEU on Dev and Test splits of AraOPUS-20 dataset. **Bolded**: best result on Test. Underlined: best result on Dev. **bg**: Bulgarian. **bs**: Bosnian. **cs**: Czech. **da**: Danish. **de**: German. **el**: Greek. **en**: English. **es**: Spanish. **fr**: French. **hi**: Hindi. **it**: Italian. **ko**: Korean. **pl**: Polish. **pt**: Portuguese. **ru**: Russian. **tr**: Turkish. **ceb**: Cebuano. **gd**: Scots Gaelic. **tmh**: Tamashek. **yo**: Yoruba. **Dash** (–): language is not supported by mBART50 many-to-many.

Model	Split	en	es	fr	ru
S2SMT	Dev	19.79	17.03	13.47	15.84
	Test	18.66	17.56	14.38	14.61
mBART	Dev	9.95	1.78	2.03	1.27
	Test	9.65	1.86	2.12	1.96
mT5	Dev	27.68	23.54	20.22	20.09
	Test	29.93	25.49	21.66	20.94
TURJ	Dev	<u>30.54</u>	<u>26.21</u>	<u>22.82</u>	<u>22.87</u>
	Test	32.07	28.16	24.11	23.95

Table 6: Results of TURJUMAN in BLEU on Dev and Test splits of UN dataset. **en**: English. **es**: Spanish. **fr**: French. **ru**: Russian.

4.4. TURJUMAN Arguments.

Each of the command-line tools (i.e., *turjuman-interactive*, *turjuman-translate*, and *turjuman-score*) support/require several arguments. Table 4 presents the required and optional arguments for each of these TURJUMAN tools.

4.5. Discussion

Results reported in Table 5 show that TURJUMAN achieves best BLEU score in 13 out of the 20 tests splits, outperforming all our baselines: S2SMT, mBART, and mT5 with +8.07, +11.28, and +0.53 BLEU points on average. We also note that mT5 outperforms AraT5 mostly in the languages that were not included in AraT5 vocabulary. Namely, we observe this in Hindi, Polish, Korean, Scots Gaelic, Tamashek, and Yoruba (see Section 2.5.). In addition, as Table 6 shows, TURJUMAN outperforms *all* baselines in UN-Test data in the four investigated MT tasks:

⁷ English, French, Spanish, and Russian → Arabic.

5. TURJUMAN: Getting Started

5.1. Installation

TURJUMAN is implemented in Python and can be installed using the pip package manager.¹⁶ It is compatible with *Python 3.6* and later versions, *Torch 1.8.1*, and *HuggingFace Transformers 4.5.1 library*.¹⁷

```
> pip install turjuman
```

5.2. Turjuman Command Line Examples

As explained, TURJUMAN provides several command-line tools for translation and evaluation. Each command supports multiple arguments. In the following, we provide a number of examples illustrating how to use TURJUMAN command-line tools with different arguments.

- **turjuman.interactive.** In the following two examples we use *turjuman.interactive* to generate translations interactively for an English and Portuguese sentences, respectively. Here, we use beam search with a beam size of 5, a maximum sequence length of 300, and a number of targets to output at 3.

```
> turjuman_interactive
```

Output

```
> Turjuman Interactive CLI
> Loading model from UBC-NLP/turjuman
> Type your source text or (q) to STOP:

> She thought a dark moment in her past
was forgotten.
> target1: اعتقدت أن لحظة مظلمة في ماضيها قد نسيت.
> target2: لقد اعتقدت أن لحظة مظلمة في ماضيها قد نسيت.
> target3: كانت تعتقد أن لحظة مظلمة في ماضيها قد نسيت.

> Type your source text or (q) to STOP:
> Esta é uma lista de estados soberanos
> target1: هذه قائمة الدول ذات السيادة
> target2: هذه قائمة بالدول ذات السيادة
> target3: قائمة الدول ذات السيادة هذه
```

- **turjuman.translate.** In the following we show how to use *turjuman.translate* with two modes of input:

(1) **Text.** A raw text is passed to the *turjuman* model directly through command line using the argument *-text* or *-t*. Translation will display directly on the terminal.

```
> turjuman_translate --text "Je peux payer
le traitement de votre fille"
```

Output

¹⁶ *pip install turjuman*

¹⁷ Installation instructions and documentation can be found at: <https://github.com/UBC-NLP/turjuman>.

```
> Turjuman Translate CLI
> Translate from input sentence
> Loading model from UBC-NLP/turjuman
> target: يمكنني أن أدفع ثمن علاج ابنتك
```

(2) **File.** The argument *-input_file* or *-f* can be used to import a set of sentences from a text file. Translation will be saved on a JSON file format.

```
> turjuman_translate --file
./sample.txt
```

Output

```
> Turjuman Translate CLI
> Translate from samples.txt
> Loading model from UBC-NLP/turjuman
> Translation is saved in samples.json
```

- **turjuman.score.** This evaluates an output translation (output translations) against reference translation(s) in terms of a BLEU score.

```
> turjuman_score -p
"translated_targets.txt" -g
"gold_targets.txt"
```

Output

```
> Turjuman Score CLI
> hyp_file=translated_targets.txt
> ref_file=gold_targets.txt
> bleu score: 43.573826221233
```

6. Conclusion

We presented TURJUMAN, an open-source Python-based package and command-line tool for Arabic neural machine translation. In the context of developing TURJUMAN, we also extracted and prepared a high quality 20 language pairs benchmark for MSA MT (*AraOPUS-20*). We exploit *AraOPUS-20* to fine-tune an Arabic text-to-text Transformer model, *AraT5*. Our resulting multilingual model outperforms competitive baselines, demonstrating the utility of our tool. In addition to its translation ability, TURJUMAN integrates a number of decoding methods. This allows for use of the tool for paraphrasing foreign sentences into diverse Arabic sequences. TURJUMAN is extensible, and we plan to train it with larger datasets and further enhance its functionality in the future. We also plan to explore adding new language pairs to TURJUMAN.

Ethical Considerations

TURJUMAN is developed using publicly available data. Hence, we do not have serious concerns about personal information being retrievable from our trained model. Similar to many NLP tools, TURJUMAN can be misused. However, the tool can be deployed for a wide host of useful application such as in education or travel. We do encourage deploying TURJUMAN in socially-relevant scenarios.

Acknowledgements

We gratefully acknowledge support from the Natural Sciences and Engineering Research Council of Canada (NSERC; RGPIN-2018-04267), the Social Sciences and Humanities Research Council of Canada (SSHRC; 435-2018-0576; 895-2020-1004; 895-2021-1008), Canadian Foundation for Innovation (CFI; 37771), Compute Canada (CC),¹⁸ UBC ARC-Sockeye,¹⁹ and Advanced Micro Devices, Inc. (AMD). Any opinions, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSERC, SSHRC, CFI, CC, AMD, or UBC ARC-Sockeye.

7. Bibliographical References

- Abdul-Mageed, M., Zhang, C., Hashemi, A., and Nagoudi, E. M. B. (2019). AraNet: A Deep Learning Toolkit for Arabic Social Media. *arXiv preprint arXiv:1912.13072*.
- Akeel, M. and Mishra, R. (2014). Ann and rule based method for english to arabic machine translation. *Int. Arab J. Inf. Technol.*, 11(4):396–405.
- Almahairi, A., Cho, K., Habash, N., and Courville, A. (2016). First Result on Arabic Neural Machine Translation. *arXiv preprint arXiv:1606.02680*.
- Alrajeh, A. (2018). A Recipe for Arabic-English Neural Machine Translation. *arXiv preprint arXiv:1808.06116*.
- Aqlan, F., Fan, X., Alqwbani, A., and Al-Mansoub, A. (2019). Arabic–chinese neural machine translation: Romanized arabic as subword unit for arabic-sourced translation. *IEEE Access*, 7:133122–133135.
- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bakr, H. A., Shaalan, K., and Ziedan, I. (2008). A hybrid approach for converting written egyptian colloquial dialect into diacritized arabic. In *The 6th international conference on informatics and systems, infos2008. cairo university*.
- Baniata, L. H., Park, S., and Park, S.-B. (2018). A neural machine translation model for arabic dialects that utilizes multitask learning (mtl). *Computational intelligence and neuroscience*, 2018.
- Bouamor, H., Habash, N., and Oflazer, K. (2014). A multidialectal parallel corpus of arabic. In *LREC*, pages 1240–1245.
- Bouamor, H., Habash, N., Salameh, M., Zaghouni, W., Rambow, O., Abdulrahim, D., Obeid, O., Khalifa, S., Eryani, F., Erdmann, A., et al. (2018). The madar arabic dialect corpus and lexicon. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- Cettolo, M., Niehues, J., Stüker, S., Bentivogli, L., and Federico, M. (2013). Report on the 10th iwslt evaluation campaign. In *Proceedings of the International Workshop on Spoken Language Translation, Heidelberg, Germany*.
- Cettolo, M., Jan, N., Sebastian, S., Bentivogli, L., Cattoni, R., and Federico, M. (2016). The iwslt 2016 evaluation campaign. In *International Workshop on Spoken Language Translation*.
- Cettolo, M., Federico, M., Bentivogli, L., Jan, N., Sebastian, S., Katsuiho, S., Koichiro, Y., and Christian, F. (2017). Overview of the iwslt 2017 evaluation campaign. In *International Workshop on Spoken Language Translation*, pages 2–14.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2009). Introduction to algorithms.
- Darwish, K. and Mubarak, H. (2016). Farasa: A new fast and accurate Arabic word segmenter. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1070–1074, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- Durrani, N., Dalvi, F., Sajjad, H., and Vogel, S. (2017). Qcri machine translation systems for iwslt 16. *arXiv preprint arXiv:1701.03924*.
- El-Kahlout, I. D., Bektaş, E., Erdem, N. Ş., and Kaya, H. (2019). Translating between morphologically rich languages: An arabic-to-turkish machine translation system. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 158–166.
- Elmahdy, M., Hasegawa-Johnson, M., and Mustafawi, E. (2014). Development of a tv broadcasts speech recognition system for qatari arabic. In *LREC*, pages 3057–3061.
- Fadaee, M., Bisazza, A., and Monz, C. (2017). Data augmentation for low-resource neural machine translation. *arXiv preprint arXiv:1705.00440*.
- Fan, A., Lewis, M., and Dauphin, Y. (2018). Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.
- Farha, I. A. and Magdy, W. (2019). Mazajak: An online arabic sentiment analyser. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 192–198.
- Farhan, W., Talafha, B., Abuammar, A., Jaikat, R., Al-Ayyoub, M., Tarakji, A. B., and Toma, A. (2020). Unsupervised dialectal neural machine translation. *Information Processing & Management*, 57(3):102181.
- Federico, M., Cettolo, M., Bentivogli, L., Michael, P., and Sebastian, S. (2012). Overview of the iwslt 2012 evaluation campaign. In *IWSLT-International Workshop on Spoken Language Translation*, pages 12–33.
- Ghoneim, M. and Diab, M. (2013). Multiword expressions in the context of statistical machine translation. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1181–1187.
- Guellil, I., Azouaou, F., and Abbas, M. (2017). Neural vs statistical translation of algerian arabic dialect written with arabizi and arabic letter. In *The 31st Pacific Asia Conference on Language, Information and Computation PACLIC*, volume 31, page 2017.
- Habash, N. and Hu, J. (2009). Improving arabic-chinese statistical machine translation using english as pivot lan-

¹⁸<https://www.computecanada.ca>

¹⁹<https://arc.ubc.ca/ubc-arc-sockeye>

- guage. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 173–181.
- Habash, N., Zalmout, N., Taji, D., Hoang, H., and Alzate, M. (2017). A parallel corpus for evaluating machine translation between Arabic and European languages. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 235–241, Valencia, Spain, April. Association for Computational Linguistics.
- Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. (2019). The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Junczys-Dowmunt, M., Dwojak, T., and Hoang, H. (2016). Is neural machine translation ready for deployment. *A case study on*, 30.
- Koehn, P. (2009). *Statistical machine translation*. Cambridge University Press.
- Lample, G., Ott, M., Conneau, A., Denoyer, L., and Ranzato, M. (2018). Phrase-based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2020). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Hoest, Q., del Moral, A. V., Jernite, Y., Thakur, A., von Platen, P., Patil, S., Chaumond, J., Drame, M., Plu, J., Tunstall, L., et al. (2021). Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184.
- Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., and Zettlemoyer, L. (2020). Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Meftouh, K., Harrat, S., Jamoussi, S., Abbas, M., and Smaili, K. (2015). Machine translation experiments on PADIC: A parallel Arabic Dialect corpus. In *Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation*, pages 26–34, Shanghai, China, October.
- Mohamed, E., Mohit, B., and Oflazer, K. (2012). Transforming standard arabic to colloquial arabic. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 176–180.
- Mubarak, H. (2018). Dial2msa: A tweets corpus for converting dialectal arabic to modern standard arabic. In *OSACT 3: The 3rd Workshop on Open-Source Arabic Corpora and Processing Tools*, page 49.
- Nagoudi, E. M. B., Elmadany, A., and Abdul-Mageed, M. (2021). Investigating code-mixed Modern Standard Arabic-Egyptian to English machine translation. In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 56–64, Online, June. Association for Computational Linguistics.
- Nagoudi, E. M. B., Elmadany, A., and Abdul-Mageed, M. (2022). Arat5: Text-to-text transformers for arabic language generation. Online, May. Association for Computational Linguistics.
- Noll, E., Oudah, M., and Habash, N. (2019). Simple automatic post-editing for arabic-japanese machine translation. *arXiv preprint arXiv:1907.06210*.
- Obeid, O., Zalmout, N., Khalifa, S., Taji, D., Oudah, M., Alhafni, B., Eryani, F., Erdmann, A., and Habash, N. (2020a). Camel tools: An open source python toolkit for arabic natural language processing. 05.
- Obeid, O., Zalmout, N., Khalifa, S., Taji, D., Oudah, M., Alhafni, B., Inoue, G., Eryani, F., Erdmann, A., and Habash, N. (2020b). CAMEL tools: An open source python toolkit for Arabic natural language processing. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 7022–7032, Marseille, France, May. European Language Resources Association.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. (2019). fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038*.
- Pasha, A., Al-Badrashiny, M., Diab, M., El Kholy, A., Eskander, R., Habash, N., Pooleery, M., Rambow, O., and Roth, R. (2014). MADAMIRA: A fast, comprehensive tool for morphological analysis and disambiguation of Arabic. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1094–1101, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2019). Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Reimers, N. and Gurevych, I. (2020). Making monolingual sentence embeddings multilingual using knowledge distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525.
- Sajjad, H., Darwish, K., and Belinkov, Y. (2013). Translating dialectal arabic to english. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–6.
- Sajjad, H., Abdelali, A., Durrani, N., and Dalvi, F. (2020). Arabench: Benchmarking dialectal arabic-english machine translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5094–5107.
- Salloum, W. and Habash, N. (2011). Dialectal to standard arabic paraphrasing to improve arabic-english statistical machine translation. In *Proceedings of the first workshop on algorithms and resources for modelling of di-*

- lects and language varieties*, pages 10–21.
- Salloum, W. and Habash, N. (2013). Dialectal arabic to english machine translation: Pivoting through modern standard arabic. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 348–358.
- Salloum, W., Elfardy, H., Alamir-Salloum, L., Habash, N., and Diab, M. (2014). Sentence level dialect identification for machine translation system selection. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 772–778.
- Song, K., Tan, X., Qin, T., Lu, J., and Liu, T.-Y. (2019). Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- Steinberger, R., Pouliquen, B., Widiger, A., Ignat, C., Erjavec, T., Tufiş, D., and Varga, D. (2006). The jrc-acquis: A multilingual aligned parallel corpus with 20+ languages. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*.
- Takezawa, T., Kikui, G., Mizushima, M., and Sumita, E. (2007). Multilingual spoken language corpus development for communication research. In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 12, Number 3, September 2007: Special Issue on Invited Papers from ISCSLP 2006*, pages 303–324.
- Tiedemann, J. (2012). Parallel data, tools and interfaces in opus. 2012:2214–2218.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A., and Raffel, C. (2020). mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Zbib, R., Malchiodi, E., Devlin, J., Stallard, D., Matsoukas, S., Schwartz, R., Makhoul, J., Zaidan, O., and Callison-Burch, C. (2012). Machine translation of arabic dialects. In *Proceedings of the 2012 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pages 49–59.
- Ziemski, M., Junczys-Dowmunt, M., and Pouliquen, B. (2016). The united nations parallel corpus v1. 0. In *Lrec*.