

ArzEn-ST: A Three-way Speech Translation Corpus for Code-Switched Egyptian Arabic - English

Injy Hamed,^{1,2} Nizar Habash,¹ Slim Abdennadher,³ Ngoc Thang Vu²

¹Computational Approaches to Modeling Language Lab, New York University Abu Dhabi

²Institute for Natural Language Processing, University of Stuttgart

³Computer Science Department, The German University in Cairo

injy.hamed@nyu.edu

Abstract

We present our work on collecting ArzEn-ST, a code-switched Egyptian Arabic - English Speech Translation Corpus. This corpus is an extension of the ArzEn speech corpus, which was collected through informal interviews with bilingual speakers. In this work, we collect translations in both directions, monolingual Egyptian Arabic and monolingual English, forming a three-way speech translation corpus. We make the translation guidelines and corpus publicly available. We also report results for baseline systems for machine translation and speech translation tasks. We believe this is a valuable resource that can motivate and facilitate further research studying the code-switching phenomenon from a linguistic perspective and can be used to train and evaluate NLP systems.

1 Introduction

Code-switching (CSW), defined as the alternation of language in text or speech, is a common linguistic phenomenon in multilingual societies. CSW can occur on the boundaries of sentences, words (within the same sentence), or morphemes (within the same word). While the worldwide prevalence of CSW has been met with increasing efforts in NLP systems trying to handle such mixed input, data sparsity remains one of the main bottlenecks hindering the development of such systems (Çetinoğlu et al., 2016).

In this paper, we present ArzEn-ST,¹ a speech translation (ST) corpus for code-switched Egyptian Arabic (Egy) - English. We extend the ArzEn Egyptian Arabic-English CSW conversational speech corpus (Hamed et al., 2020) with translations going to both directions; the primary (Egyptian Arabic) as well as secondary (English) languages. See Figure 1. This corpus is a valuable resource filling an important gap, given the naturalness and high

¹Arz is the ISO 639-3 code for Egyptian Arabic.

Audio:



Transcription:

بحس ان ده بيعمل نتيجة عكسية يعني

Egy Translation:

بحس ان ده بيعمل نتيجة عكسية يعني

English Translation:

I feel that's counter-productive actually

Figure 1: An example from the corpus, showing the four representations for each utterance: audio, transcription, Egyptian Arabic translation, and English translation.

frequency of CSW in it. It can be used for the purpose of linguistic investigations as well as for building and evaluating NLP systems. We provide benchmark baseline results for the tasks of automatic speech recognition (ASR), machine translation (MT), and ST. We make the translation guidelines and full corpus available, as well as the experiments' scripts and data splits.²

The paper is organized as follows. In Section 2, we provide an overview of previous work done for code-switched ASR, MT, and ST tasks as well as corpora collection. In Section 3, we provide an overview of the ArzEn speech corpus. In Section 4, we elaborate on the translation guidelines used to create the three-way parallel ST corpus. Finally, in Section 5, we report the performance of the ASR, MT, and ST baseline systems.

2 Related Work

2.1 CSW Automatic Speech Recognition

CSW ASR has gained a considerable amount of research (Vu et al., 2012; Li and Vu, 2019; Ali et al., 2021; Hamed et al., 2022a; Hussein et al., 2022), where several CSW speech corpora have been collected, covering multiple language pairs, including Chinese-English (Lyu et al., 2015), Hindi-English (Ramanarayanan and Suendermann-Oeft, 2017), Spanish-English (Solorio and Liu, 2008), Arabic-

²<http://arzen.camel-lab.com/>

| Language | Citations |
|--------------------------------------|---|
| Translating CSW → monolingual | |
| Hindi-English → English | (Dhar et al., 2018; Srivastava and Singh, 2020; Tarunesh et al., 2021; Chen et al., 2022) |
| Sinhala-English → Sinhala | (Kugathasan and Sumathipala, 2021) |
| English-Spanish → English | (Chen et al., 2022) |
| MSA-Egyptian Arabic → English | (Chen et al., 2022) |
| English-Bengali → both | (Mahata et al., 2019) |
| Egyptian Arabic -English → both | ArzEn-ST (the corpus presented in this paper) |
| Translating monolingual → CSW | |
| Hindi → Hindi-English | (Tarunesh et al., 2021; Banerjee et al., 2018) |
| English → Bengali-English | (Banerjee et al., 2018) |
| English → Gujarati-English | (Banerjee et al., 2018) |
| English → Tamil-English | (Banerjee et al., 2018) |
| English, Hindi → Hindi-English | (Srivastava and Singh, 2021) |

Table 1: Overview on available human-annotated CSW-focused parallel corpora.

English (Ismail, 2015; Hamed et al., 2018, 2020; Chowdhury et al., 2021), Arabic-French (Djegdjiga et al., 2018), Frisian-Dutch (Yilmaz et al., 2016), Mandarin-Taiwanese (Lyu et al., 2006), Turkish-German (Çetinoğlu, 2017), English-Malay (Ahmed and Tan, 2012), English-isiZulu (van der Westhuizen and Niesler, 2016) and Sepedi-English (Modipa et al., 2013). In this work, we build on our ArzEn speech corpus (Hamed et al., 2020), and enrich it with multiple translations.

2.2 CSW Machine Translation

While research in CSW MT has been gaining attention over the past years (Sinha and Thakur, 2005; Dhar et al., 2018; Mahata et al., 2019; Menacer et al., 2019; Song et al., 2019; Tarunesh et al., 2021; Xu and Yvon, 2021; Chen et al., 2022; Hamed et al., 2022b; Gaser et al., 2022), the collected CSW parallel corpora are limited. By looking into the reported corpora, we identify a number of dimensions in which they vary. First is synthetic or human-annotated data. Second, for human-annotated data, it can be either collected, or especially commissioned for MT/NLP. Third, for collected data, it can be obtained from textual or speech sources. And finally, the data set may include translations to one or more languages.

To circumvent the data scarcity issue, researchers investigated the use of synthetically generated CSW parallel data for training and testing (Gupta et al., 2020; Yang et al., 2020; Xu and Yvon, 2021). While this is acceptable for training purposes, synthetic data should not be used for testing,

as it does not reflect real-world CSW distributions.

For collecting human-annotated parallel corpora, researchers have tried either asking bilingual speakers to translate naturally occurring CSW sentences into monolingual sentences, or as another solution to data scarcity, have commissioned annotators to translate monolingual sentences into CSW sentences. In Table 1, we present a summary of available human-annotated parallel corpora that are focused on CSW. For the latter approach, we note that generating CSW data in a human-commissioned fashion could differ from naturally-occurring CSW sentences. Such data could be biased to the grammatical structure of the monolingual sentences, and could be dominated by single noun switches, being the easiest CSW type to generate.

The former approach of translating naturally-occurring CSW sentences into monolingual sentences is the most optimal way to collect a CSW parallel corpus; however, most of the collected corpora rely on CSW sentences obtained from textual sources (mostly from social media platforms). The main concern here is that CSW phenomena occurring in text are more restricted than those occurring in natural speech. In text, people are usually dissuaded from changing scripts, and therefore either avoid switching languages, or switch languages without switching scripts. The latter issue was tackled by Shazal et al. (2020), where the authors used a sequence-to-sequence deep learning model to transliterate SMS/chat text collected by Chen et al. (2017) from Arabizi (where Arabic words

are written in Roman script) to Arabic orthography. While this corpus is not focused on CSW, it contains CSW sentences.

Finally, we categorize the collected corpora in terms of the translation direction. As shown in Table 1, most of the corpora include translations for CSW sentences to the secondary language, which is most commonly English. A smaller number of researchers investigated translating CSW sentences into the primary language. And even fewer researchers included translations to both directions.

The work of Menacer et al. (2019) is also relevant to our work. The authors extracted Modern Standard Arabic (MSA)-English CSW sentences from the UN documents, to which English translations are available (Eisele et al., 2010). The Arabic translations were generated by translating the English segments using the Google Translate API. While this can be used for training purposes, these translations should not be used as gold reference. Moreover, given the nature of the corpus, it contained limited types of CSW, as opposed to the types that occur in conversational speech.

2.3 CSW Speech Translation

Work on CSW ST is still in its early stages, with little prior work (Nakayama et al., 2019; Weller et al., 2022; Huber et al., 2022). For CSW ST corpora, two corpora are available for Spanglish: Bangor Miami (Cieri et al., 2004) and Fisher (Deuchar et al., 2014). While the Fisher dataset is not a CSW-focused corpus, it contains a considerable amount of CSW (Weller et al., 2022). Similarly, for CSW Egyptian Arabic-English, the Callhome dataset also contains some amount of CSW (Gadalla et al., 1997; Kumar et al., 2014). A Japanese-English ST corpus (Nakayama et al., 2019) was also collected, however it includes read-speech and not spontaneous speech. Huber et al. (2022) collected a one-hour German-English code-switching speech translation corpus containing read-speech.

Our new corpus, ArzEn-ST, fills an important resource gap, providing an ST corpus for code-switched Egyptian Arabic-English. The corpus is human-annotated where the source sentences are collected through interviews with bilingual speakers, and contain naturally-generated CSW sentences; they are then translated in both directions: monolingual Egyptian Arabic and monolingual English.

3 Overview of the ArzEn Corpus

ArzEn is a conversational speech corpus that is collected through informal interviews. The interviews were held at the German University in Cairo, which is a private university where English is the instruction language. The topics discussed were general topics such as education, work and life experiences, career, technology, personal life, hobbies, and travelling experiences. No instructions were given to participants regarding code-switching; they were not asked to produce nor avoid code-switching. Interviews were held with 38 Egyptian Arabic-English bilingual speakers (61.5% males, 38.5% females), in the age range of 18-35, who are students (55%) and employees (45%) at the university. The speech corpus comprises of 12 hours of speech, containing 6,216 sentences.

3.1 Code-switching Types in ArzEn

The four main CSW types mentioned in the literature are present in ArzEn (Poplack, 1980; Stefanich et al., 2019). We present a corpus example for each of the types in Table 2.

Inter-sentential CSW This type of CSW is defined as switching languages from one sentence to another.

Extra-sentential CSW This type of CSW, also called **tag-switching**, is where tag elements from one language are inserted into a monolingual sentence in another language, without the need for grammatical considerations. It mostly involves the use of fillers, interjections, tags, and idiomatic expressions. This type of CSW requires only minimal knowledge of the grammar of the secondary language.

Intra-sentential CSW This type, also referred to as **code-mixing**, is defined as using multiple languages within the same sentence, where the CSW segments must conform to the underlying syntactic rules of both languages. This type of CSW requires a better understanding of the grammar of both languages, compared to extra-sentential CSW.

Intra-word CSW This type, also called **morphological CSW**, is where switching occurs at the level of morphemes. Given that Egyptian Arabic is a morphologically rich language (Habash et al., 2012b), morphological code-switching occurs where Egyptians attach Arabic clitics and affixes to English words.

| CSW Type | Example |
|----------------------|---|
| Inter-sentential CSW | It's very difficult making friends at work. عملت صحاب بس مش فى الشغل. <i>I made friends, but not at work.</i> It's very difficult making friends at work. |
| Extra-sentential CSW | أنا مولود فى مصر فى القاهرة Okay Okay <i>I was born in Egypt, in Cairo.</i> |
| Intra-sentential CSW | كان فى جزء actually related to research <i>There was actually a part</i> related to research |
| Intra-word CSW | ايه ال expectation+بتاعك لل project+؟ <i>What is your expectation for the project?</i> |
| Explicatory CSW | Okay طيب فى مثلا quote أو مقولة قريتها فى كتاب وعجبتك؟ Okay <i>okay is there a quote or a quote</i> that you read in a book and liked? |
| Elaboratory CSW | لو ال novel دي يعني هي oriental او هي شرقية فقرأها اه من .. من author عربي <i>If this novel is like oriental or it is [originally written] in Arabic then I will read it ah from .. from an Arab author [in its Arabic version].</i> |

Table 2: Examples of different CSW types in ArzEn followed by their English translation. The originally Arabic phrases are italicized in the English translation. For Explicatory and Elaboratory CSW, the underlining marks the repeated phrases.

In addition to the above, and motivated by our interest in translation from CSW texts, we identify two types of repetitive CSW phenomena in terms of their communicative purposes.

Explicatory CSW This type of CSW is where the speaker simply repeats the same word in another language.

Elaboratory CSW This type of CSW is where the speaker code-switches to further elaborate on the meaning.

Both types are challenging in terms of handling the CSW repetitions when translating into a single language. We address these issues in Section 4.

3.2 Code-switching Statistics in ArzEn

ArzEn contains a considerable amount of CSW. On the sentence level, 33.2% of the sentences are monolingual Arabic, 3.1% are monolingual English, and 63.7% code-mixed. Among the code-mixed sentences, 46.0% have morphological CSW.

On the word level, in the code-mixed sentences, 81.3% of the words are Arabic, 15.2% are English, and 3.4% are morphologically CSW words. Morphological CSW in ArzEn involves the use of both Arabic clitics and affixes. A list of the clitics and affixes occurring in morphological code-switched words present in the ArzEn corpus and their frequencies are provided in Hamed et al. (2022a).

3.3 Input Transcription

The ArzEn collected interviews were manually transcribed by Egyptian Arabic-English bilingual speakers. The transcribers were requested to use Arabic script for Arabic words and Roman script for English words. For morphological CSW words, Arabic clitics and affixes are written in Arabic script and English words are written in Roman script, as follows: **Arabic prefixes/proclitics + English words # Arabic suffixes/enclitics**, for example *ت #TASK+ ال Al+TASK#A³ 'the+task#s'*. While the transcribers generally followed the rules in a strict manner, we observe script confusion in the case of borrowed words that have become strongly embedded in Egyptian Arabic. In such cases, transcriptions can contain occurrences of the same words in both scripts, such as *mobile* and *موبايل mwbAyl, film* and *فيلم fylm*, and *camera* and *كاميرا kAmyrA*.

Given the spontaneous nature of the corpus, disfluencies were found due to repetitions, corrections, and changing course/structure mid-sentence. Such disfluencies were marked with ‘..’, which occurs in more than 26% of the corpus sentences. The following tags were also used for non-speech parts: [HES] for hesitation, [HUM] for humming, [COUGH], [LAUGHTER], and [NOISE].

³Transliteration in the HSB scheme (Habash et al., 2007).

4 ArzEnST Translation Guidelines

The transcriptions are translated to monolingual English and monolingual Egyptian Arabic sentences by human translators.⁴ In this section, we discuss the translation guidelines. In general, our decisions are mainly guided by giving a higher priority to fluency over accuracy. We opt for producing as natural as possible outputs that reflect the style of the original sentence. Even though we acknowledge that some of our decisions can make the translation task harder for MT systems, our goal is to produce natural translations. The guidelines cover three categories, general translation rules (denoted by *GR*), conversational speech translation rules (denoted by *SR*), and code-switching translation rules (denoted by *CSWR*). In Table 3, we present translation examples covering some of the rules.

4.1 General Translation Rules (*GR*)

[*GR_{intended}*] Translators are requested to provide natural translations with the intended meaning rather than literal translations. This also covers the case of idiomatic expressions. See Table 3 (a).

[*GR_{difficult}*] Similar to the LDC Arabic-to-English Translation Guidelines (LDC, 2013), segments that are difficult to translate should be indicated using ((text)). Such cases usually contain highly dialectal Arabic words or Arabic idioms. See Table 3 (b).

[*GR_{abbrev}*] For all abbreviations, we made the decision to provide transliteration as pronounced instead of translation,⁵ for example *NLP* is transliterated as ان ال بي *An Al by*, and *AIESEC* is transliterated as آيزيك *Āzyk*.

[*GR_{propn}*] Non-abbreviated proper nouns should be transliterated, unless they have meaning. In that case, they should be translated as long as the meaning of the sentence remains coherent, otherwise, should be transliterated.⁶ See Table 3 (c).

⁴English translations are performed by one translator, and the dev and test sets are revised by one of the authors. Egyptian Arabic translations are performed by one translator and revised by another.

⁵We plan to annotate these cases with full translations in the future, to assist in tasks interested in removing English/Arabic text.

⁶The translators were advised to refer to Wikipedia Arabic for the translations of titles of books and films.

4.2 Conversational Speech Translation Rules (*SR*)

[*SR_{style}*] Translations should capture the same fluency and style of the original text. This means that disfluencies such as repetitions should also be included in translations. See Table 3 (d).

[*SR_{punc+}*] Punctuation, non-speech tags, and disfluency marks ‘.’ present in the source text should be kept the same and in the same relative position in the sentence in the target translation.

[*SR_{partial}*] Due to disfluencies, it is common to have partial Arabic words. We transliterate such partial words, and similar to LDC (2013), we mark them with a preceding ‘%’ sign. See Table 3 (e).

4.3 Code-switching Translation Rules (*CSWR*)

[*CSWR_{borrowed}*] For English words that are commonly used in Arabic, an attempt should first be made to identify a commonly used reasonable translation, otherwise, translators are allowed to transliterate. Examples of the latter case, included loanwords such as *mobile* and *laptop* that have become strongly integrated in Arabic, as opposed to *online* and *presentation* which can be translated to عرض تقديمي and عبر الأنترنت, respectively.

[*CSWR_{rewrites}*] We allow modifications to CSW segments when translating into English for the purpose of achieving better fluency. Similarly, when translating into Arabic, we also allow slight modifications to the original Arabic words. We elaborate on such cases for both directions below.

CSW→En: We allow modification to the original English words. This is mainly needed to handle difference in grammatical structures across languages as well as morphological CSW. For example, ASK+ بي *by+ASK* is translated as ‘he asks’. See Table 3 (f-g).

CSW→Ar: It is allowed to slightly modify the original Arabic words for better fluency. The following are common cases where this is needed.

- Since Arabic makes heavy use of the definite article + ال *Al+* to mark different constructions such as adjectival modification and idafa (possessive construct), translators are given permission to drop/reassign the placement of definite articles for the purpose of maximizing fluency. For example, the adjectival construction

| | | |
|---|---|---|
| General Translation Rules | (a) | [GRintended] Provide intended meaning rather than literal translation |
| | CSW: | لأنه ال+team ممكن أخش مع حد مابيشتعلمش فأشيل أنا الليلة كلها |
| | Egy: | لأنه الفريق ممكن أخش مع حد مابيشتعلمش فأشيل أنا الليلة كلها |
| Eng: | Because in the team, I can be with someone that doesn't work and ((I will be responsible for everything)). | |
| General Translation Rules | (b) | [GRdifficult] Indicate segments that are hard to translate |
| | CSW: | فكست علي ال+internship و روح ال+camp معاهم |
| | Egy: | فكست علي التدريب و روح المعسكر معاهم |
| Eng: | ((So, I declined)) the internship and I went to the camp with them. | |
| General Translation Rules | (c) | [GRpropn] Translate or transliterate proper nouns |
| | CSW: | أكثر فيلم بحبه The Dark Knight, Batman's Dark Knight عشان بحس إن يعني الفيلم ده فيه كمية أفكار عبقرية مش طبيعية |
| | Egy: | أكثر فيلم بحبه فارس الظلام, باتمان فارس الظلام عشان بحس إن يعني الفيلم ده فيه كمية أفكار عبقرية مش طبيعية |
| Eng: | The movie I like the most is The Dark Knight, Batman's Dark Knight because I feel that this movie has a lot of creative ideas. | |
| Conversational Speech Translation Rules | (d) | [SRstyle] Capture the same meaning, fluency, and naturalness |
| | CSW: | سافرت برا مصر قبل كده غير ال .. ال+bachelor project؟ |
| | Egy: | سافرت برا مصر قبل كده غير ال .. مشروع التخرج؟ |
| Eng: | Have you ever travelled abroad apart from the .. the bachelor project? | |
| Conversational Speech Translation Rules | (e) | [SRpartial] Transliterate and annotate partial words |
| | CSW: | طيب و إزاي بت ..بت+overcome الموضوع ده؟ إزاي بتعدي ال .. ال+stress ده؟ |
| | Egy: | طيب و إزاي بت ..بتتخطى الموضوع ده؟ إزاي بتعدي ال .. الضغط ده؟ |
| Eng: | Well, how do you %bt .. #overcome this issue?how do you overpass this .. this #stress? | |
| Code-switching Translation Rules | (f) | [CSWRewrites] Modify English words in translation |
| | CSW: | [HUM] و قعدت في كذا city |
| | Egy: | [HUM] و قعدت في كذا مدينة |
| | Eng: | [HUM] and I have stayed in multiple cities |
| | (g) | [CSWRewrites] Modify English words in translation |
| | CSW: | فانا لو عندك كذا robot و كذا task المفروض يتعملوا, فانا بكتب ال+code لل robots إن هي تروح لل+task ات دي, بس |
| | Egy: | فانا لو عندك كذا روبوت و كذا مهمة المفروض يتعملوا, فانا بكتب الكود للروبوتات إن هي تروح للمهام دي, بس |
| | Eng: | So, if you have multiple robots and multiple tasks you should be working on, so I just write the code for the robots so that they should handle these tasks , that's it. |
| | (h) | [CSWRewrites] Modify Arabic words in translation |
| | CSW: | يعني ببقي environment مختلف عن ال+environment بتاع ال+big city |
| Egy: | يعني بتبقى بيئة مختلفة عن البيئة بتاعة المدن الكبيرة | |
| Eng: | I mean the environment differs from the environment of the big city. | |
| (i) | [CSWRewrites] Modify Arabic words in translation | |
| CSW: | [HES] هو ال+bachelor project بتاعى هو يعني related somehow ال+optical field | |
| Egy: | [HES] هو مشروع التخرج بتاعى هو يعني بشكل ما مرتبط بالمجال البصري | |
| Eng: | [HES] my bachelor project is actually somehow related to the optical field | |
| (j) | [CSRWreorder] Modify the order between Arabic and English words | |
| CSW: | لكن لو عندنا stereo camera في طرق معينة نقدر نستخدمها إن إحنا نقدر نحدد المسافة بين الكاميرا أو بين ال+ two objects معينين بالطريقة دي تمام؟ | |
| Egy: | لكن لو عندنا كاميرا ستيريو في طرق معينة نقدر نستخدمها إن إحنا نقدر نحدد المسافة بين الكاميرا أو بين جسمين معينين بالطريقة دي تمام؟ | |
| Eng: | But if we have a stereo camera and in certain ways we can use it to estimate the distance between the camera and the specific two objects this way, okay? | |
| (k) | [CSRWstyle] Handle repetitions in segments with syntactic divergence | |
| CSW: | ماشى okay طيب ممكن تحكيلنا عن ال .. ال+bachelor project بتاعك ؟ | |
| Egy: | ماشى تمام طيب ممكن تحكيلنا عن ال .. مشروع التخرج بتاعك؟ | |
| Eng: | Ok,well,can you tell us about your..your bachelor project ? | |
| CSW: | اه [HES] أكثر .. من أكثر ال .. my favorite books يعني animal .. animal farm | |
| Egy: | اه [HES] أكثر .. من أكثر ال .. كتيبي المفضلة يعني مزرعة .. مزرعة الحيوان | |
| Eng: | Yes [HES] the most .. one of the most .. my favourite books I mean .. animal animal farm . | |

Table 3: Translation examples following different guideline rules.

working life + *ال* gets translated to *الحياة المهنية* (with two definite article instances).

- Given that Arabic is a gender-marking language, gender reinflection is sometimes needed to guarantee fluent translation into Arabic. See Table 3 (h).
- Modifying Arabic prepositions following English words. See Table 3 (i).

[*CSWR_{reorder}*] We allow changing the order between the original Arabic and English words to handle syntactic divergences between the two languages and achieve better fluency. See Table 3 (j).

[*CSWR_{style}*] The *SR_{style}* rule gets further compounded in the case of CSW when repetition occurs at a location where there is syntactic divergence between both languages, such as adjectival phrases in our language pair. In this case, since the order of words changes during translation, the word to be repeated at that position could also change. Following our fluency preference, we prefer the translation that gives higher fluency over providing an accurate literal translation. See Table 3 (k).

[*CSWR_{disfluency}*] Another interesting translation challenge arises in the context of CSW when speakers repeat words using different languages. In the case of explicatory CSW, where the English and Arabic words have the exact same meaning, we allow translating the English word into the same present Arabic word, treating it as a case of repetition due to disfluency. In the case of elaboratory CSW, where CSW is used to further elaborate on meaning, we ask the translator to find another translation of the word that would better capture the subtle difference between both words. If that is not possible, we allow the repetition of the word.

5 Benchmarking Baseline Systems

In this section, we discuss the ASR, MT, and ST baseline systems. We describe the experimental setup for each and present the results in Table 4.

5.1 Experimental Setup

We follow the same train, dev, and test splits defined in Hamed et al. (2020). For all the experiments, we use ArzEn-ST dev set (1,402 sentences) for tuning and ArzEn-ST test set (1,470 sentences) for testing. For training, we use ArzEn-ST train set (3,344 sentences), in addition to other monolingual data which we mention below.

Automatic Speech Recognition We train a joint CTC/attention based E2E ASR system using ES-Pnet (Watanabe et al., 2018). The encoder and decoder consist of 12 and 6 Transformer blocks with 4 heads, feed-forward inner dimension 2048 and attention dimension 256. The CTC/attention weight (λ_1) is set to 0.3. SpecAugment (Park et al., 2019) is applied for data augmentation. For the Language Model (LM), the RNNLM consists of 1 LSTM layer with 1000 hidden units and is trained for 20 epochs. For decoding, the beam size is 20 and the CTC weight is 0.2.

In addition to using ArzEn-ST for training, we also train the ASR system and LM using Callhome (Gadalla et al., 1997), MGB-3 (Ali et al., 2017), a 5-hours subset from Librispeech (Panayotov et al., 2015), and a 5-hours subset from MGB-2 (Ali et al., 2016).⁷ We perform Alif/Ya normalization (Arabic), remove punctuation and corpus-specific annotations, and lower-case English words.⁸

Machine Translation We train Transformer models using Fairseq (Ott et al., 2019) on a single GeForce RTX 3090 GPU. We use the hyperparameters from the FLORES benchmark for low-resource machine translation (Guzmán et al., 2019). The hyperparameters are given in Appendix A. For each MT model, we use a BPE model trained jointly on source and target sides. The BPE model is trained using Fairseq with `character_coverage` set to 1.0. We tune the vocabulary size for each experiment for the values of $1k$, $3k$, $5k$, $8k$, and $16k$.

In addition to ArzEn-ST, we also train the MT system using $324k$ extra Egyptian Arabic-English parallel sentences obtained from the following parallel corpora: Callhome Egyptian Arabic-English Speech Translation Corpus (Kumar et al., 2014), LDC2012T09 (Zbib et al., 2012), LDC2017T07 (Chen et al., 2017), LDC2019T01 (Chen et al., 2019), LDC2020T05 (Li et al., 2020), and MADAR (Bouamor et al., 2018).⁹ These extra corpora include $15k$ sentences with CSW instances. When translating into En, we use all these extra corpora as Arabic-English training. However, when translating into Egy, we use these extra corpora as English-Arabic training, but we exclude the $15k$ sentences producing CSW Arabic as our reference does not have CSW sentences. Data pre-

⁷We followed the setup used in (Hamed et al., 2022a).

⁸For the Callhome corpus, we removed partial words.

⁹For corpora with no defined data splits, we follow the guidelines provided in Diab et al. (2013).

| Training Set | ASR | | MT | | ST | |
|------------------|------|------|--------|---------|--------|---------|
| | WER | CER | CSW→En | CSW→Egy | CSW→En | CSW→Egy |
| | | | BLEU | BLEU | BLEU | BLEU |
| ArzEn-ST | 57.9 | 36.2 | 8.6 | 48.0 | 4.5 | 13.0 |
| ArzEn-ST + Extra | 34.7 | 20.0 | 34.3 | 79.8 | 16.5 | 31.1 |

Table 4: Summary of results for baseline systems evaluated on ArzEn-ST test set. We present baseline systems for both settings: (1) training using ArzEn-ST data only and (2) training using ArzEn-ST data with **Extra** monolingual speech corpora for the ASR system and **Extra** monolingual Egyptian Arabic-English parallel sentences for the MT systems. We report Word Error Rate (WER) and Character Error Rate (CER) for ASR systems, and BLEU score (Papineni et al., 2002) using SacrebleuBLEU (Post, 2018) for MT and ST systems.

processing involved removing all corpus-specific annotations, URLs and emoticons, lowercasing, running Moses’ (Koehn et al., 2007) tokenizer, MADAMIRA (Pasha et al., 2014) simple tokenization (D0) and Alif/Ya normalization (Arabic).

Speech Translation We build a cascaded speech translation system, where we train an ASR system and use an MT system to translate the ASR system’s outputs. We opt for a cascaded system over an end-to-end system due to the limitation of available resources to build an end-to-end system, in addition to the fact that cascaded systems have been shown to outperform end-to-end systems in low-resource settings (Denisov et al., 2021).

5.2 Results

Table 4 presents the results for the MT and ST baseline systems. We also report results for the ASR system used to build the cascaded ST system.¹⁰ We report results for both settings: (1) when training only using ArzEn-ST corpus and (2) when training using ArzEn-ST corpus in addition to the extra monolingual data specified for each task (**Extra**). As expected, adding extra monolingual data greatly improves results. We observe that translating into Arabic achieves higher BLEU scores than translating into English. This is expected, as in the case of translating from CSW text, Arabic words (around 85% of words) remain mostly the same with possible slight modifications required. We also observe that for the ST models, the performance is nearly reduced by half compared to the MT results. This highlights the difficulty of the task. Given that CSW ST has only been slightly tackled by other researchers, we hope that this corpus will motivate further research on this task.

¹⁰ASR results are different than those reported in Hamed et al. (2022a) as we limit the data to publicly-available corpora, use different preprocessing, and different data splits.

6 Conclusion and Future Work

Code-switching has become a worldwide prevalent phenomenon. This created a need for NLP systems to be able to handle such mixed input. Code-switched data is typically scarce, which is evident in the limited number of available corpora for machine translation and speech translation tasks. In this paper, we extend the previously collected ArzEn speech corpus with translations to both its primary and secondary languages, providing a three-way code-switched Egyptian Arabic-English speech translation corpus. We have discussed the translation guidelines, particularly with regards to issues arising due to the spontaneous nature of the corpus as well as code switching. We reported benchmark results for baseline ASR, MT, and ST systems. We make this corpus available to motivate and facilitate further research in this area.

For future work, we plan on improving the corpus and using it for code-switching linguistic investigations as well as NLP tasks. With regards to corpus improvements, we plan on adding additional translation references and CODAfyng (Habash et al., 2012a; Eskander et al., 2013) the corpus. From a linguistic perspective, having signals from the monolingual Arabic and English translations, we plan to further understand why code-switching occurs at the given points. Finally, we plan to use this corpus for NLP tasks, working on data augmentation for the purpose of improving machine translation and speech translation.

Acknowledgements

This project has benefited from financial support by DAAD (German Academic Exchange Service). We also thank the reviewers for their insightful comments and constructive feedback.

References

- Basem HA Ahmed and Tien-Ping Tan. 2012. Automatic speech recognition of code switching speech using 1-best rescoring. In *Proceedings of the International Conference on Asian Language Processing*, pages 137–140.
- Ahmed Ali, Peter Bell, James Glass, Yacine Messaoui, Hamdy Mubarak, Steve Renals, and Yifan Zhang. 2016. The mgb-2 challenge: Arabic multi-dialect broadcast media recognition. In *Proceedings of the Spoken Language Technology Workshop*, pages 279–284.
- Ahmed Ali, Shammur Chowdhury, Amir Hussein, and Yasser Hifny. 2021. Arabic code-switching speech recognition using monolingual data. In *Proceedings of Interspeech*.
- Ahmed Ali, Stephan Vogel, and Steve Renals. 2017. Speech recognition challenge in the wild: Arabic mgb-3. In *Proceedings of the Automatic Speech Recognition and Understanding Workshop*, pages 316–322.
- Suman Banerjee, Nikita Moghe, Siddhartha Arora, and Mitesh M Khapra. 2018. A dataset for building code-mixed goal oriented conversation systems. In *Proceedings of the International Conference on Computational Linguistics*.
- Houda Bouamor, Nizar Habash, Mohammad Salameh, Wajdi Zaghouni, Owen Rambow, Dana Abdulrahim, Ossama Obeid, Salam Khalifa, Fadhl Eryani, Alexander Erdmann, and Kemal Oflazer. 2018. The MADAR Arabic dialect corpus and lexicon. In *Proceedings of the International Conference on Language Resources and Evaluation*.
- Özlem Çetinoğlu. 2017. A code-switching corpus of turkish-german conversations. In *Proceedings of the Linguistic Annotation Workshop*, pages 34–40.
- Özlem Çetinoğlu, Sarah Schulz, and Ngoc Thang Vu. 2016. Challenges of computational processing of code-switching. In *Proceedings of the Workshop on Computational Approaches to Linguistic Code Switching*.
- Shuguang Chen, Gustavo Aguilar, Anirudh Srinivasan, Mona Diab, and Tamar Solorio. 2022. Calcs 2021 shared task: Machine translation for code-switched data. *arXiv preprint arXiv:2202.09625*.
- Song Chen, Dana Fore, Stephanie Strassel, Haejoong Lee, and Jonathan Wright. 2017. BOLT Egyptian Arabic sms/chat and transliteration LDC2017T07. Philadelphia: Linguistic Data Consortium.
- Song Chen, Jennifer Tracey, Christopher Walker, and Stephanie Strassel. 2019. BOLT Egyptian Arabic parallel discussion forums data. Linguistic Data Consortium (LDC) catalog number LDC2019T01, ISBN 1-58563-871-4.
- Shammur Absar Chowdhury, Amir Hussein, Ahmed Abdelali, and Ahmed Ali. 2021. Towards one model to rule all: Multilingual strategy for dialectal code-switching Arabic asr. In *Proceedings of Interspeech*.
- Christopher Cieri, David Miller, and Kevin Walker. 2004. The fisher corpus: A resource for the next generations of speech-to-text. In *Proceedings of the International Conference on Language Resources and Evaluation*, pages 69–71.
- Pavel Denisov, Manuel Mager, and Ngoc Thang Vu. 2021. IMS’systems for the IWSLT 2021 low-resource speech translation task. *Proceedings of the International Conference on Spoken Language Translation*.
- Margaret Deuchar, Peredur Davies, Jon Herring, M Carmen Parafita Couto, and Diana Carter. 2014. Building bilingual corpora. *Advances in the Study of Bilingualism*, pages 93–111.
- Mrinal Dhar, Vaibhav Kumar, and Manish Shrivastava. 2018. Enabling code-mixed translation: Parallel corpus creation and MT augmentation approach. In *Proceedings of the Workshop on Linguistic Resources for Natural Language Processing*, pages 131–140.
- Mona Diab, Nizar Habash, Owen Rambow, and Ryan Roth. 2013. LDC Arabic treebanks and associated corpora: Data divisions manual. *arXiv preprint arXiv:1309.5652*.
- Amazouz Djegdjiga, Martine Adda-Decker, and Lori Lamel. 2018. The French-Algerian code-switching triggered audio corpus (FACST). In *Proceedings of the International Conference on Language Resources and Evaluation*.
- Andreas Eisele, Yu Chen, and UN Multi. 2010. a multilingual corpus from united nation documents. In *Proceedings of the International Conference on Language Resources and Evaluation*, pages 924–929.
- Ramy Eskander, Nizar Habash, Owen Rambow, and Nadi Tomeh. 2013. Processing spontaneous orthography. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 585–595, Atlanta, Georgia.
- Hassan Gadalla, Hanaa Kilany, Howaida Arram, Ashraf Yacoub, Alaa El-Habashi, Amr Shalaby, Krisjanis Karins, Everett Rowson, Robert MacIntyre, Paul Kingsbury, David Graff, and Cynthia McLemore. 1997. Callhome Egyptian Arabic transcripts. *Linguistic Data Consortium, Philadelphia*.
- Marwa Gaser, Manuel Mager, Injy Hamed, Nizar Habash, Slim Abdennadher, and Ngoc Thang Vu. 2022. Exploring segmentation approaches for neural machine translation of code-switched Egyptian Arabic-English text. *arXiv preprint arXiv:2210.06990*.
- Deepak Gupta, Asif Ekbal, and Pushpak Bhattacharyya. 2020. A semi-supervised approach to generate the code-mixed text using pre-trained encoder and transfer learning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2267–2280.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. The flores evaluation datasets for low-resource machine

- translation: Nepali-English and Sinhala-English. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Nizar Habash, Mona Diab, and Owen Rambow. 2012a. Conventional orthography for dialectal Arabic. In *Proceedings of the International Conference on Language Resources and Evaluation*, pages 711–718.
- Nizar Habash, Ramy Eskander, and Abdelati Hawwari. 2012b. A morphological analyzer for Egyptian Arabic. In *Proceedings of the meeting of the special interest group on computational morphology and phonology*, pages 1–9.
- Nizar Habash, Abdelhadi Soudi, and Tim Buckwalter. 2007. On Arabic Transliteration. In A. van den Bosch and A. Soudi, editors, *Arabic Computational Morphology: Knowledge-based and Empirical Methods*, pages 15–22. Springer, Netherlands.
- Injy Hamed, Pavel Denisov, Chia-Yu Li, Mohamed Elmahdy, Slim Abdennadher, and Ngoc Thang Vu. 2022a. Investigations on speech recognition systems for low-resource dialectal Arabic-English code-switching speech. *Computer Speech & Language*, 72:101278.
- Injy Hamed, Mohamed Elmahdy, and Slim Abdennadher. 2018. Collection and analysis of code-switch Egyptian Arabic-English speech corpus. In *Proceedings of the International Conference on Language Resources and Evaluation*.
- Injy Hamed, Nizar Habash, Slim Abdennadher, and Ngoc Thang Vu. 2022b. Investigating lexical replacements for Arabic-English code-switched data augmentation. *arXiv preprint arXiv:2205.12649*.
- Injy Hamed, Ngoc Thang Vu, and Slim Abdennadher. 2020. Arzen: A speech corpus for code-switched Egyptian Arabic-English. In *Proceedings of the International Conference on Language Resources and Evaluation*, pages 4237–4246.
- Christian Huber, Enes Yavuz Ugan, and Alexander Waibel. 2022. Code-switching without switching: Language agnostic end-to-end speech translation. *arXiv preprint arXiv:2210.01512*.
- Amir Hussein, Shammur Absar Chowdhury, Ahmed Abdelali, Najim Dehak, and Ahmed Ali. 2022. Code-switching text augmentation for multilingual speech processing. *arXiv preprint arXiv:2201.02550*.
- Manal A Ismail. 2015. The sociolinguistic dimensions of code-switching between Arabic and English by Saudis. *International Journal of English Linguistics*, 5(5):99.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics Demo and Poster Sessions*, pages 177–180.
- Archchana Kugathanan and Sagara Sumathipala. 2021. Neural machine translation for Sinhala-English code-mixed text. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, pages 718–726.
- Gaurav Kumar, Yuan Cao, Ryan Cotterell, Chris Callison-Burch, Daniel Povey, and Sanjeev Khudanpur. 2014. Translations of the callhome Egyptian Arabic corpus for conversational speech translation. In *Proceedings of International Workshop on Spoken Language Translation*.
- Linguistic Data Consortium LDC. 2013. [BOLT program: Arabic to English translation guidelines](#).
- Chia-Yu Li and Ngoc Thang Vu. 2019. Integrating knowledge in end-to-end automatic speech recognition for Mandarin-English code-switching. In *International Conference on Asian Language Processing*, pages 160–165.
- Xuansong Li, Stephen Grimes, and Stephanie Strassel. 2020. BOLT Egyptian Arabic-English word alignment – conversational telephone speech training. Linguistic Data Consortium (LDC) catalog number LDC2020T05, ISBN 1-58563-920-6.
- Dau-Cheng Lyu, Ren-Yuan Lyu, Yuang-chin Chiang, and Chun-Nan Hsu. 2006. Speech recognition on code-switching among the Chinese dialects. In *Proceedings of the International Conference on Acoustics Speech and Signal Processing Proceedings*, volume 1, pages I–I.
- Dau-Cheng Lyu, Tien-Ping Tan, Eng-Siong Chng, and Haizhou Li. 2015. Mandarin-English code-switching speech corpus in south-east asia: Seame. *Language Resources and Evaluation*, 49(3):581–600.
- Sainik Kumar Mahata, Soumil Mandal, Dipankar Das, and Sivaji Bandyopadhyay. 2019. Code-mixed to monolingual translation framework. In *Proceedings of the Forum for Information Retrieval Evaluation*, pages 30–35.
- Mohamed Amine Menacer, David Langlois, Denis Jovet, Dominique Fohr, Odile Mella, and Kamel Smaïli. 2019. Machine translation on a parallel code-switched corpus. In *Proceedings of the Canadian Conference on Artificial Intelligence*, pages 426–432. Springer.
- Thipe I Modipa, Marelise H Davel, and Febe De Wet. 2013. Implications of Sepedi/English code switching for ASR systems. In *Proceedings of Pattern Recognition Association of South Africa*.
- Sahoko Nakayama, Takatomo Kano, Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2019. Recognition and translation of code-switching speech utterances. In *Proceedings of the Conference of the Oriental COCODA International Committee for the Coordination and Standardisation of Speech Databases and Assessment Techniques*, pages 1–6.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. Fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*.

- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an ASR corpus based on public domain audio books. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, pages 5206–5210.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. 2019. Specaugment: A simple data augmentation method for automatic speech recognition. In *Proceeding of Interspeech*, pages 2613–2617.
- Arfath Pasha, Mohamed Al-Badrashiny, Mona T Diab, Ahmed El Kholy, Ramy Eskander, Nizar Habash, Manoj Pooleery, Owen Rambow, and Ryan Roth. 2014. Madamira: A fast, comprehensive tool for morphological analysis and disambiguation of Arabic. In *Proceedings of the International Conference on Language Resources and Evaluation*, pages 1094–1101.
- Shana Poplack. 1980. Sometimes i’ll start a sentence in Spanish y termino en espanol: toward a typology of code-switching. *Linguistics*, 18:581–618.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Conference on Machine Translation: Research Papers*, pages 186–191.
- Vikram Ramanarayanan and David Suendermann-Oeft. 2017. Jee haan, i’d like both, por favor: Elicitation of a code-switched corpus of Hindi-English and Spanish-English human-machine dialog. In *Proceedings of Interspeech*, pages 47–51.
- Ali Shazal, Aiza Usman, and Nizar Habash. 2020. A unified model for Arabizi detection and transliteration using sequence-to-sequence models. In *Proceedings of the Arabic Natural Language Processing Workshop*, pages 167–177.
- R Mahesh K Sinha and Anil Thakur. 2005. Machine translation of bi-lingual Hindi-English (Hinglish) text. In *Proceedings of the Machine Translation summit (MT Summit X)*, pages 149–156.
- Tamar Solorio and Yang Liu. 2008. Learning to predict code-switching points. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 973–981.
- Kai Song, Yue Zhang, Heng Yu, Weihua Luo, Kun Wang, and Min Zhang. 2019. Code-switching for enhancing NMT with pre-specified translation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 449–459.
- Vivek Srivastava and Mayank Singh. 2020. Phinc: A parallel Hinglish social media code-mixed corpus for machine translation. In *Proceedings of the Workshop on Noisy User-generated Text*.
- Vivek Srivastava and Mayank Singh. 2021. Hinge: A dataset for generation and evaluation of code-mixed Hinglish text. In *Proceedings of the Workshop on Evaluation and Comparison of NLP Systems*.
- Sara Stefanich, Jennifer Cabrelli, Dustin Hilderman, and John Archibald. 2019. The morphophonology of intraword codeswitching: Representation and processing. *Frontiers in Communication*, page 54.
- Ishan Tarunesh, Syamantak Kumar, and Preethi Jyothi. 2021. From machine translation to code-switching: Generating high-quality code-switched text. In *Proceedings of The Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*.
- Ewald van der Westhuizen and Thomas Niesler. 2016. Automatic speech recognition of English-isiZulu code-switched speech from South African soap operas. *Procedia Computer Science*, 81:121–127.
- Ngoc Thang Vu, Dau-Cheng Lyu, Jochen Weiner, Dominic Telaar, Tim Schlippe, Fabian Blaicher, Eng-Siong Chng, Tanja Schultz, and Haizhou Li. 2012. A first speech recognition system for Mandarin-English code-switch conversational speech. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, pages 4889–4892. IEEE.
- Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, et al. 2018. Espnet: End-to-end speech processing toolkit. In *Proceeding of Interspeech*, pages 2207–2207.
- Orion Weller, Matthias Sperber, Telmo Pires, Hendra Setiawan, Christian Gollan, Dominic Telaar, and Matthias Paulik. 2022. End-to-end speech translation for code switched speech. *Findings of the Association for Computational Linguistics*.
- Jitao Xu and François Yvon. 2021. Can you traduir this? machine translation for code-switched input. In *Proceedings of the Workshop on Computational Approaches to Linguistic Code-Switching*, pages 84–94.
- Zhen Yang, Bojie Hu, Ambyera Han, Shen Huang, and Qi Ju. 2020. Csp: Code-switching pre-training for neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2624–2636.
- Emre Yilmaz, Maaïke Andringa, Sigrid Kingma, Jelske Dijkstra, F Kuip, H Velde, Frederik Kampstra, Jouke Algra, H Heuvel, and David A van Leeuwen. 2016. A longitudinal bilingual Frisian-Dutch radio broadcast database designed for code-switching research. In *Proceedings of the International Conference on Language Resources and Evaluation*.
- Rabih Zbib, Erika Malchiodi, Jacob Devlin, David Stalard, Spyros Matsoukas, Richard Schwartz, John Makhoul, Omar F. Zaidan, and Chris CallisonBurch. 2012. Machine translation of Arabic dialects. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pages 49–59.

A MT Hyperparameters

The following is the train command:

```
python3 fairseq_cli/train.py $DATA_DIR --source-lang src --target-lang tgt --arch transformer --share-all-embeddings --encoder-layers 5 --decoder-layers 5 --encoder-embed-dim 512 --decoder-embed-dim 512 --encoder-ffn-embed-dim 2048 --decoder-ffn-embed-dim 2048 --encoder-attention-heads 2 --decoder-attention-heads 2 --encoder-normalize-before --decoder-normalize-before --dropout 0.4 --attention-dropout 0.2 --relu-dropout 0.2 --weight-decay 0.0001 --label-smoothing 0.2 --criterion label_smoothed_cross_entropy --optimizer adam --adam-betas '(0.9, 0.98)' --clip-norm 0 --lr-scheduler inverse_sqrt --warmup-updates 4000 --warmup-init-lr 1e-7 --lr 1e-3 --stop-min-lr 1e-9 --max-tokens 4000 --update-freq 4 --max-epoch 100 --save-interval 10 --ddp-backend=no_c10d
```