

Code-Switched Language Models Using Neural Based Synthetic Data from Parallel Sentences

Genta Indra Winata, Andrea Madotto, Chien-Sheng Wu, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

{giwinata, amadotto, cwuak}@connect.ust.hk, pascale@ece.ust.hk

Abstract

Training code-switched language models is difficult due to lack of data and complexity in the grammatical structure. Linguistic constraint theories have been used for decades to generate artificial code-switching sentences to cope with this issue. However, this requires external word alignments or constituency parsers that create erroneous results on distant languages. We propose a sequence-to-sequence model using a copy mechanism to generate code-switching data by leveraging parallel monolingual translations from a limited source of code-switching data. The model learns how to combine words from parallel sentences and identifies when to switch one language to the other. Moreover, it captures code-switching constraints by attending and aligning the words in inputs, without requiring any external knowledge. Based on experimental results, the language model trained with the generated sentences achieves state-of-the-art performance and improves end-to-end automatic speech recognition.

1 Introduction

Code-switching is a common linguistic phenomenon in multilingual communities, in which a person begins speaking or writing in one language and then switches to another in the same sentence.¹ It is motivated in response to social factors as a way of communicating in a multicultural society. In its practice, code-switching varies due to the traditions, beliefs, and normative values in the respective communities. Linguists have studied the code-switching phenomenon and proposed a number of linguistic theories (Poplack, 1978; Pfaff, 1979; Poplack, 1980; Belazi et al., 1994). Code-switching is not produced indiscriminately,

¹Code-switching refers to mixing of languages following the definitions in Poplack (1980). We use “intra-sentential code-switching” interchangeably with “code-mixing”.

but follows syntactic constraints. Many linguists have formulated various constraints to define a general rule for code-switching (Poplack, 1978, 1980; Belazi et al., 1994). However, these constraints cannot be postulated as a universal rule for all code-switching scenarios, especially for languages that are syntactically divergent (Berk-Seligson, 1986), such as English and Mandarin since they have word alignments with an inverted order.

Building a language model (LM) and an automatic speech recognition (ASR) system that can handle intra-sentential code-switching is known to be a difficult research challenge. The main reason lies in the unpredictability of code-switching points in an utterance and data scarcity. Creating a large-scale code-switching dataset is also very expensive. Therefore, code-switching data generation methods to augment existing datasets are a useful workaround.

Existing methods that apply equivalence constraint theory to generate code-switching sentences (Li and Fung, 2012; Pratapa et al., 2018) may suffer performance issues as they receive erroneous results from the word aligner and the part-of-speech (POS) tagger. Thus, this approach is not reliable and effective. Recently, Garg et al. (2018) proposed a SeqGAN-based model to generate code-switching sentences. Indeed, the model learns how to generate new synthetic sentences. However, the distribution of the generated sentences is very different from real code-switching data, which leads to underperforming results.

To overcome the challenges in the existing works, we introduce a neural-based code-switching data generator model using pointer-generator networks (Pointer-Gen) (See et al., 2017) to learn code-switching constraints from a limited source of code-switching data and leverage their translations in both languages. Intu-

itively, the copy mechanism can be formulated as an end-to-end solution to copy words from parallel monolingual sentences by aligning and reordering the word positions to form a grammatical code-switching sentence. This method solves the two issues in the existing works by removing the dependence on the aligner or tagger, and generating new sentences with a similar distribution to the original dataset. Interestingly, this method can learn the alignment effectively without a word aligner or tagger. As an additional advantage, we demonstrate its interpretability by showing the attention weights learned by the model that represent the code-switching constraints. Our contributions are summarized as follows:

- We propose a language-agnostic method to generate code-switching sentences using a pointer-generator network (See et al., 2017) that learns when to switch and copy words from parallel sentences, without using external word alignments or constituency parsers. By using the generated data in the language model training, we achieve the state-of-the-art performance in perplexity and also improve the end-to-end ASR on an English-Mandarin code-switching dataset.
- We present an implementation applying the equivalence constraint theory to languages that have significantly different grammar structures, such as English and Mandarin, for sentence generation. We also show the effectiveness of our neural-based approach in generating new code-switching sentences compared to the equivalence constraint and SeqGAN (Garg et al., 2018).
- We thoroughly analyze our generation results and further examine how our model identifies code-switching points to show its interpretability.

2 Generating Code-Switching Data

In this section, we describe our proposed model to generate code-switching sentences using a pointer-generator network. Then, we briefly list the assumptions of the equivalence constraint (EC) theory, and explain our application of EC theory for sentence generation. We call the dominant language the matrix language (L_1) and the inserted language the embedded language (L_2), following

the definitions from Myers-Scotton (2001). Let us define $Q = \{Q_1, \dots, Q_T\}$ as a set of L_1 sentences and $E = \{E_1, \dots, E_T\}$ as a set of L_2 sentences with T number of sentences, where each $Q_t = \{q_{1,t}, \dots, q_{m,t}\}$ and $E_t = \{e_{1,t}, \dots, e_{n,t}\}$ are sentences with m and n words. E is the corresponding parallel sentences of Q .

2.1 Pointer-Gen

Initially, Pointer-Gen was proposed to learn when to copy words directly from the input to the output in text summarization, and they have since been successfully applied to other natural language processing tasks, such as comment generation (Lin et al., 2019). The Pointer-Gen leverages the information from the input to ensure high-quality generation, especially when the output sequence consists of elements from the input sequence, such as code-switching sequences.

We propose to use Pointer-Gen by leveraging parallel monolingual sentences to generate code-switching sentences. The approach is depicted in Figure 1. The pointer-generator model is trained from concatenated sequences of parallel sentences (Q, E) to generate code-switching sentences, constrained by code-switching texts. The words of the input are fed into the encoder. We use a bidirectional long short-term memory (LSTM), which, produces hidden state h_t in each step t . The decoder is a unidirectional LSTM receiving the word embedding of the previous word. For each decoding step, a generation probability $p_{gen} \in [0, 1]$ is calculated, which weights the probability of generating words from the vocabulary, and copying words from the source text. p_{gen} is a soft gating probability to decide whether to generate the next token from the decoder or to copy the word from the input instead. The attention distribution a_t is a standard attention with general scoring (Luong et al., 2015). It considers all encoder hidden states to derive the context vector. The vocabulary distribution $P_{voc}(w)$ is calculated by concatenating the decoder state s_t and the context vector h_t^* :

$$p_{gen} = \sigma(w_{h^*}^T h_t^* + w_s^T s_t + w_x^T x_t + b_{ptr}), \quad (1)$$

where w_{h^*} , w_s , and w_x are trainable parameters and b_{ptr} is the scalar bias. The vocabulary distribution $P_{voc}(w)$ and the attention distribution a^t are weighted and summed to obtain the final distribution $P(w)$, which is calculated as follows:

$$P(w) = p_{gen} P_{voc}(w) + (1 - p_{gen}) \sum_{i:w_i=w} a_i^t. \quad (2)$$

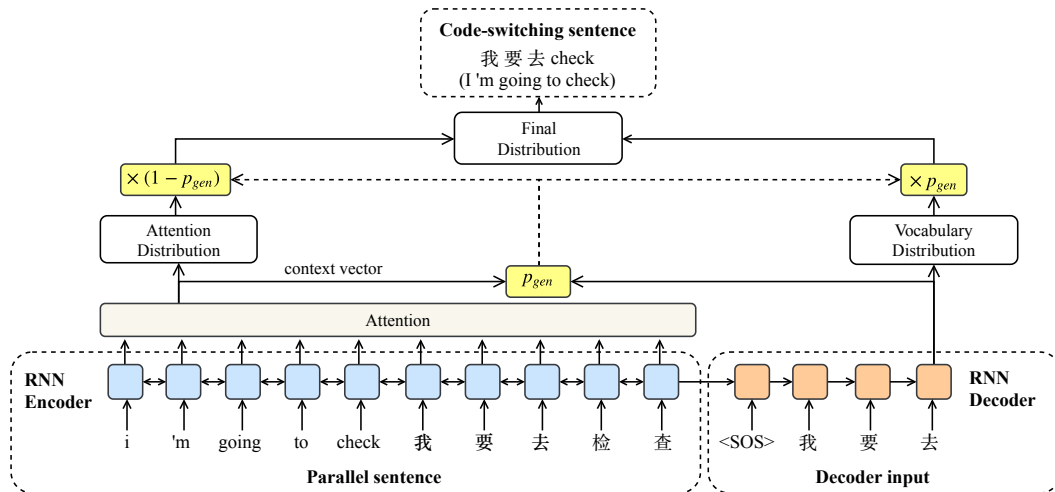


Figure 1: **Pointer-Gen** model, which includes an RNN encoder and RNN decoder. The parallel sentence is the input of the encoder, and in each decoding step, the decoder generates a new token.

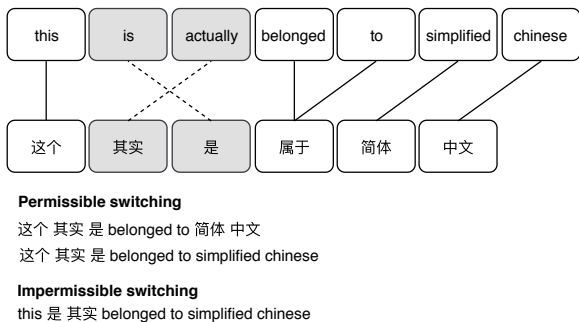


Figure 2: Example of equivalence constraint (Li and Fung, 2012). Solid lines show the alignment between the matrix language (top) and the embedded language (bottom). The dotted lines denote impermissible switching.

We use a beam search to select the N -best code-switching sentences.

2.2 Equivalence Constraint

Studies on the EC (Poplack, 1980, 2013) show that code-switching only occurs where it does not violate the syntactic rules of either language. An example of a English-Mandarin mixed-language sentence generation is shown in Figure 2, where EC theory does not allow the word “其实” to come after “是” in Chinese, or the word “is” to come after “actually”. Pratapa et al. (2018) apply the EC in English-Spanish language modeling with a strong assumption. We are working with English and Mandarin, which have distinctive grammar structures (e.g., part-of-speech tags), so applying a constituency parser would give us erroneous results. Thus, we simplify sentences into a linear structure, and we allow lexical substitution on non-

crossing alignments between parallel sentences. Alignments between an L_1 sentence Q_t and an L_2 sentence E_t comprise a source vector with indices $u_t = \{a_1, a_2, \dots, a_m\} \in \mathbb{W}^m$ that has a corresponding target vector $v_t = \{b_1, b_2, \dots, b_m\} \in \mathbb{W}^m$, where u is a sorted vector of indices in an ascending order. The alignment between a_i and b_i does not satisfy the constraint if there exists a pair of a_j and b_j , where $(a_i < a_j, \text{ and } b_i > b_j)$ or $(a_i > a_j, \text{ and } b_i < b_j)$. If the switch occurs at this point, it changes the grammatical order in both languages; thus, this switch is not acceptable. During the generation step, we allow any switches that do not violate the constraint. We propose to generate synthetic code-switching data by the following steps:

1. Align the L_1 sentences Q and L_2 sentences E using `fast_align`² (Dyer et al., 2013). We use the mapping from the L_1 sentences to the L_2 sentences.
2. Permute alignments from step (1) and use them to generate new sequences by replacing the phrase in the L_1 sentence with the aligned phrase in the L_2 sentence.
3. Evaluate generated sequences from step (2) if they satisfy the EC theory.

3 End-to-End Code-Switching ASR

To show the effectiveness of our proposed method, we build a transformer-based end-to-end code-

²The code implementation can be found at https://github.com/clab/fast_align.

switching ASR system. The end-to-end ASR model accepts a spectrogram as the input, instead of log-Mel filterbank features (Zhou et al., 2018), and predicts characters. It consists of N layers of an encoder and decoder. Convolutional layers are added to learn a universal audio representation and generate input embedding. We employ multi-head attention to allow the model to jointly attend to information from different representation subspaces at a different position.

For proficiency in recognizing individual languages, we train a multilingual ASR system trained from monolingual speech. The idea is to use it as a pretrained model and transfer the information while training the model with code-switching speech. This is an effective method to initialize the parameters of low-resource ASR such as code-switching. The catastrophic forgetting issue arises when we train one language after the other. Therefore, we solve the issue by applying a multi-task learning strategy. We jointly train speech from both languages by taking the same number of samples for each language in every batch to keep the information of both tasks.

In the inference time, we use beam search, selecting the best sub-sequence scored using the softmax probability of the characters. We define $P(Y)$ as the probability of the sentence. We incorporate language model probability $p_{lm}(Y)$ to select more natural code-switching sequences from generation candidates. A word count is added to avoid generating very short sentences. $P(Y)$ is calculated as follows:

$$P(Y) = \alpha P_{trans}(Y|X) + \beta p_{lm}(Y) + \gamma \sqrt{wc(Y)} \quad (3)$$

where α is the parameter to control the decoding probability from the probability of characters from the decoder $P_{trans}(Y|X)$, β is the parameter to control the language model probability $p_{lm}(Y)$, and γ is the parameter to control the effect of the word count $wc(Y)$.

4 Experiments

4.1 Data Preparation

We use speech data from SEAME Phase II, a conversational English-Mandarin Chinese code-switching speech corpus that consists of spontaneously spoken interviews and conversations (Nanyang Technological University, 2015). We split the corpus following information

from Winata et al. (2018a). The details are depicted in Table 1. We tokenize words using the Stanford NLP toolkit (Manning et al., 2014). For monolingual speech datasets, we use HKUST (Liu et al., 2006), comprising spontaneous Mandarin Chinese telephone speech recordings, and Common Voice, an open-accented English dataset collected by Mozilla.³ We split Chinese words into characters to avoid word boundary issues, similarly to Garg et al. (2018). We generate L_1 sentences and L_2 sentences by translating the training set of SEAME Phase II into English and Chinese using the Google NMT system (To enable reproduction of the results, we release the translated data).⁴ Then, we use them to generate 270,531 new pieces of code-switching data, which is thrice the number of the training set. Table 2 shows the statistics of the new generated sentences. To calculate the complexity of our real and generated code-switching corpora, we use the following measures:

Switch-Point Fraction (SPF) This measure calculates the number of switch-points in a sentence divided by the total number of word boundaries (Pratapa et al., 2018). We define "switch-point" as a point within the sentence at which the languages of words on either side are different.

Code Mixing Index (CMI) This measure counts the number of switches in a corpus (Gambäck and Das, 2014). At the utterance level, it can be computed by finding the most frequent language in the utterance and then counting the frequency of the words belonging to all other languages present. We compute this metric at the corpus level by averaging the values for all the sentences in a corpus. The computation is shown as follows:

$$C_u(x) = \frac{N(x) - \max(\ell_i \in \ell\{t_{\ell_i}(x)\}) + P(x)}{N(x)}, \quad (4)$$

where $N(x)$ is the number of tokens of utterance x , t_{ℓ_i} is the tokens in language ℓ_i , and $P(x)$ is the number of code-switching points in utterance x . We compute this metric at the corpus-level by averaging the values for all sentences.

³The dataset is available at <https://voice.mozilla.org/>.

⁴We have attached the translated data in the Supplementary Materials.

4.2 LM Training Strategy Comparison

We generate code-switching sentences using three methods: EC theory, SeqGAN (Garg et al., 2018), and Pointer-Gen. To find the best way of leveraging the generated data, we compare different training strategies as follows:

- (1) rCS, (2a) EC, (2b) SeqGAN,
- (2c) Pointer-Gen, (3a) EC & rCS,
- (3b) SeqGAN & rCS, (3c) Pointer-Gen & rCS
- (4a) EC → rCS (4b) SeqGAN → rCS,
- (4c) Pointer-Gen → rCS

(1) is the baseline, training with real code-switching data. (2a–2c) train with only augmented data. (3a–3c) train with the concatenation of augmented data with rCS. (4a–4c) run a two-step training, first training the model only with augmented data and then fine-tuning with rCS. Our early hypothesis is that the results from (2a) and (2b) will not be as good as the baseline, but when we combine them, they will outperform the baseline. We expect the result of (2c) to be on par with (1), since Pointer-Gen learns patterns from the rCS dataset, and generates sequences with similar code-switching points.

4.3 Experimental Setup

In this section, we present the settings we use to generate code-switching data, and train our language model and end-to-end ASR.

Pointer-Gen The pointer-generator model has 500-dimensional hidden states. We use 50k words as our vocabulary for the source and target. We optimize the training by Stochastic Gradient Descent with an initial learning rate of 1.0 and decay of 0.5. We generate the three best sequences using beam search with five beams, and sample 270,531 sentences, thrice the amount of the code-switched training data.

EC We generate 270,531 sentences, thrice the amount of the code-switched training data. To make a fair comparison, we limit the number of switches to two for each sentence to get a similar number of code-switches (SPF and CMI) to Pointer-Gen.

SeqGAN We implement the SeqGAN model using a PyTorch implementation⁵, and use our best

⁵To implement SeqGAN, we use code from <https://github.com/suragnair/seqGAN>.

	Train	Dev	Test
# Speakers	138	8	8
# Duration (hr)	100.58	5.56	5.25
# Utterances	90,177	5,722	4,654
# Tokens	1.2M	65K	60K
CMI	0.18	0.22	0.19
SPF	0.15	0.19	0.17

Table 1: Data statistics of SEAME Phase II. The dataset is split by speaker ID.

	EC	SeqGAN	Pointer-Gen
# Utterances	270,531	270,531	270,531
# Words	3,040,202	2,981,078	2,922,941
new unigram	13.63%	34.67%	4.67%
new bigram	69.43%	80.33%	46.57%
new trigram	99.73%	141.56%	69.38%
new four-gram	121.04%	182.89%	85.07%
CMI	0.25	0.13	0.25
SPF	0.17	0.2	0.17

Table 2: Statistics of the generated data. The table shows the number of utterances and words, code-switches ratio, and percentage of new n-grams.

trained LM baseline as the generator in SeqGAN. We sample 270,531 sentences from the generator, thrice the amount of the code-switched training data (with a maximum sentence length of 20).

LM In this work, we focus on sentence generation, so we evaluate our data with the same two-layer LSTM LM for comparison. It is trained using a two-layer LSTM with a hidden size of 200 and unrolled for 35 steps. The embedding size is equal to the LSTM hidden size for weight tying (Press and Wolf, 2017). We optimize our model using SGD with an initial learning rate of 20. If there is no improvement during the evaluation, we reduce the learning rate by a factor of 0.75. In each step, we apply a dropout to both the embedding layer and recurrent network. The gradient is clipped to a maximum of 0.25. We optimize the validation loss and apply an early stopping procedure after five iterations without any improvements. In the fine-tuning step of training strategies (4a–4c), the initial learning rate is set to 1.

End-to-end ASR We convert the inputs into normalized frame-wise spectrograms from 16-kHz audio. Our transformer model consists of two encoder and decoder layers. An Adam optimizer and Noam warmup are used for training with an initial learning rate of 1e-4. The model has a hidden size of 1024, a key dimension of 64,

Training Strategy	Overall		en-zh		zh-en		en-en		zh-zh	
	valid	test	valid	test	valid	test	valid	test	valid	test
<i>Only real code-switching data</i>										
(1) rCS	72.89	65.71	7411.42	7857.75	120.41	130.21	29.31	29.61	244.88	246.71
<i>Only generated data</i>										
(2a) EC	115.98	96.54	32865.62	30580.89	107.22	109.10	28.24	28.2	1893.77	1971.68
(2b) SeqGAN	252.86	215.17	33719	37119.9	174.2	187.5	91.07	88	1799.74	1783.71
(2c) Pointer-Gen	72.78	64.67	7055.59	7473.68	119.56	133.39	27.77	27.67	234.16	235.34
<i>Concatenate generated data with real code-switching data</i>										
(3a) EC & rCS	70.33	62.43	8955.79	9093.01	130.92	139.06	26.49	26.28	227.57	242.30
(3b) SeqGAN & rCS	77.37	69.58	8477.44	9350.73	134.27	143.41	30.64	30.81	260.89	264.28
(3c) Pointer-Gen & rCS	68.49	61.57	7146.08	7667.82	127.50	139.06	26.75	26.96	218.27	226.60
<i>Pretrain with generated data and fine-tune with real code-switching data</i>										
(4a) EC → rCS	68.46	61.42	8200.78	8517.29	101.15	107.77	25.49	25.78	247.3	258.95
(4b) SeqGAN → rCS	70.61	64.03	6950.02	7694.2	114.82	122.84	28.5	28.73	236.94	244.62
(4c) Pointer-Gen → rCS	66.08	59.74	6620.76	7172.42	114.53	127.12	26.36	26.40	216.02	222.49

Table 3: Results of perplexity (PPL) on a valid set and test set for different training strategies. We report the overall PPL, and code-switching points (**en-zh**) and (**zh-en**), as well as the monolingual segments PPL (**en-en**) and (**zh-zh**).




Figure 3: Results of perplexity (PPL) on different numbers of generated samples. The graph shows that Pointer-Gen attains a close performance to the real training data, and outperforms *SeqGAN* and *EC*.

and a value dimension of 64. The training data are randomly shuffled every epoch. Our character set is the concatenation of English letters, Chinese characters found in the corpus, spaces, and apostrophes. In the multilingual ASR pretraining, we train the model for 18 epochs. Since the sizes of the datasets are different, we over-sample the smaller dataset. The fine-tuning step takes place after the pretraining using code-switching data. In the inference time, we explore the hypothesis using beam search with eight beams and a batch size of 1.

4.4 Evaluation Metrics

We employ the following metrics to measure the performance of our models.

Token-level Perplexity (PPL) For the LM, we calculate the PPL of characters in Mandarin Chinese and words in English. The reason is that some Chinese words inside the SEAME corpus are not well tokenized, and tokenization results are

not consistent. Using characters instead of words in Chinese can alleviate word boundary issues. The PPL is calculated by taking the exponential of the sum of losses. To show the effectiveness of our approach in calculating the probability of the switching, we split the perplexity computation into monolingual segments (**en-en**) and (**zh-zh**), and code-switching segments (**en-zh**) and (**zh-en**).

Character Error Rate (CER) For our ASR, we compute the overall CER and also show the individual CERs for Mandarin Chinese (**zh**) and English (**en**). The metric calculates the distance of two sequences as the *Levenshtein Distance*.

5 Results & Discussion

LM In Table 3, we can see the perplexities of the test set evaluated on different training strategies. Pointer-Gen consistently performs better than state-of-the-art models such as *EC* and *SeqGAN*. Comparing the results of models trained using only generated samples, (**2a-2b**) leads to




Figure 4: The visualization of pointer-generator attention weights on input words in each time-step during the inference time. The y-axis indicates the generated sequence, and the x-axis indicates the word input. In this figure, we show the code-switching points when our model attends to words in the L_1 and L_2 sentences: **left**: (“no”, “没有”) and (“then”, “然后”), **right**: (“we”, “我们”), (“share”, “一起”) and (“room”, “房间”).

the undesirable results that are also mentioned by Pratapa et al. (2018), but it does not apply to Pointer-Gen (2c). We can achieve a similar results with the model trained using only real code-switching data, rCS. This demonstrates the quality of our data generated using Pointer-Gen. In general, combining any generated samples with real code-switching data improves the language model performance for both code-switching segments and monolingual segments. Applying concatenation is less effective than the two-step training strategy. Moreover, applying the two-step training strategy achieves the state-of-the-art performance.

As shown in Table 2, we generate new n-grams including code-switching phrases. This leads us to a more robust model, trained with both generated data and real code-switching data. We can see clearly that Pointer-Gen-generated samples have a distribution more similar to the real code-switching data compared with SeqGAN, which shows the advantage of our proposed method.

Effect of Data Size To understand the importance of data size, we train our model with different amounts of generated data. Figure 3 shows the PPL of the models with different amounts of generated data. An interesting finding is that our model trained with only 78K samples of Pointer-Gen data (same number of samples as rCS) achieves a similar PPL to the model trained with only rCS, while SeqGAN and EC have a significantly higher PPL. We can also see that 10K samples of Pointer-Gen data is as good as 270K samples of EC data. In general, the number of samples is positively correlated with the improvement in performance.

Model	Overall	en	zh
Baseline	34.40%	41.79%	35.94%
+ Pre-training	32.76%	40.06%	32.44%
+ LM (rCS)	32.25%	39.45%	31.90%
+ LM (Pointer-Gen → rCS)	31.07%	38.39%	30.85%

Table 4: ASR evaluation, showing the performance on all sequences (Overall), English segments (en), and Mandarin Chinese segments (zh).

ASR Evaluation We evaluate our proposed sentence generation method on an end-to-end ASR system. Table 4 shows the CER of our ASR systems, as well as the individual CER on each language. Based on the experimental results, pre-training is able to reduce the error rate by 1.64%, as it corrects the spelling mistakes in the prediction. After we add LM (rCS) to the decoding step, the error rate can be reduced to 32.25%. Finally, we replace the LM with LM (Pointer-Gen → rCS), and it further decreases the error rate by 1.18%.

Model Interpretability We can interpret a Pointer-Gen model by extracting its attention matrices and then analyzing the activation scores. We show the visualization of the attention weights in Figure 4. The square in the heatmap corresponds to the attention score of an input word. In each time-step, the attention scores are used to select words to be generated. As we can observe in the figure, in some cases, our model attends to words that are translations of each other, for example, the words (“no”, “没有”), (“then”, “然后”), (“we”, “我们”), (“share”, “一起”), and (“room”, “房间”). This indicates the model can identify code-switching points, word alignments, and translations without being given any explicit information.

rCS		Pointer-Gen		
POS tags	ratio	POS tags	ratio	examples
English				
NN (noun)	56.16%	NN (noun)	55.45%	那个 consumer 是不 (that consumer is not)
RB (adverb)	10.34%	RB (adverb)	10.14%	okay so 其实 (okay so its real)
JJ (adjective)	7.04%	JJ (adjective)	7.16%	我很 jealous 的每次 (i am very jealous every time)
VB (verb)	5.88%	VB (verb)	5.89%	compared 这个 (compared to this)
Chinese				
VV (other verbs)	23.77%	VV (other verbs)	23.72%	讲的要用 用 microsoft word (i want to use microsoft word)
M (measure word)	16.83%	M (measure word)	16.49%	我们有这个 个 god of war (we have this god of war)
DEG (associative)	9.12%	DEG (associative)	9.13%	我们 的 result (our result)
NN (common noun)	9.08%	NN (common noun)	8.93%	我应该不会讲 话 because intimidated by another (i shouldn't talk because intimidated by another)

Table 5: The most common English and Mandarin Chinese part-of-speech tags that trigger code-switching. We report the frequency ratio from **Pointer-Gen**-generated sentences compared to the real code-switching data. We also provide an example for each POS tag.

Code-Switching Patterns Table 5 shows the most common English and Mandarin Chinese POS tags that trigger code-switching. The distribution of word triggers in the Pointer-Gen data are similar to the real code-switching data, indicating our model’s ability to learn similar code-switching points. Nouns are the most frequent English word triggers. They are used to construct an optimal interaction by using cognate words and to avoid confusion. Also, English adverbs such as “then” and “so” are phrase or sentence connectors between two language phrases for intra-sentential and inter-sentential code-switching. On the other hand, Chinese transitional words such as the measure word “个” or associative word “的” are frequently used as inter-lingual word associations.

6 Related Work

Code-switching language modeling research has been focused on building a model that handles mixed-language sentences and on generating synthetic data to solve the data scarcity issue. The first statistical approach using a linguistic theory was introduced by Li and Fung (2012), who adapted the EC on monolingual sentence pairs during the decoding step of an ASR system. Ying and Fung (2014) implemented a functional-head constraint lattice parser with a weighted finite-state transducer to reduce the search space on a code-switching ASR system. Then, Adel et al. (2013a)

extended recurrent neural networks (RNNs) by adding POS information to the input layer and a factorized output layer with a language identifier. The factorized RNNs were also combined with an n-gram backoff model using linear interpolation (Adel et al., 2013b), and syntactic and semantic features were added to them (Adel et al., 2015). Baheti et al. (2017) adapted an effective curriculum learning by training a network with monolingual corpora of two languages, and subsequently trained on code-switched data. A further investigation of EC and curriculum learning showed an improvement in English-Spanish language modeling (Pratapa et al., 2018), and a multi-task learning approach was introduced to train the syntax representation of languages by constraining the language generator (Winata et al., 2018a). Garg et al. (2018) proposed to use SeqGAN (Yu et al., 2017) for generating new mixed-language sequences. Winata et al. (2018b) leveraged character representations to address out-of-vocabulary words in the code-switching named entity recognition. Finally, Winata et al. (2019) proposed a method to represent code-switching sentence using language-agnostic meta-representations.

7 Conclusion

We propose a novel method for generating synthetic code-switching sentences using Pointer-Gen by learning how to copy words from parallel cor-

pora. Our model can learn code-switching points by attending to input words and aligning the parallel words, without requiring any word alignments or constituency parsers. More importantly, it can be effectively used for languages that are syntactically different, such as English and Mandarin Chinese. Our language model trained using outperforms equivalence constraint theory-based models. We also show that the learned language model can be used to improve the performance of an end-to-end automatic speech recognition system.

Acknowledgments

This work has been partially funded by ITF/319/16FP and MRP/055/18 of the Innovation Technology Commission, the Hong Kong SAR Government, and School of Engineering Ph.D. Fellowship Award, the Hong Kong University of Science and Technology, and RDC 1718050-0 of EMOS.AI. We sincerely thank the three anonymous reviewers for their insightful comments on our paper.

References

- Heike Adel, Ngoc Thang Vu, Katrin Kirchhoff, Dominic Telaar, and Tanja Schultz. 2015. Syntactic and semantic features for code-switching factored language models. *IEEE Transactions on Audio, Speech, and Language Processing*, 23(3):431–440.
- Heike Adel, Ngoc Thang Vu, Franziska Kraus, Tim Schlippe, Haizhou Li, and Tanja Schultz. 2013a. Recurrent neural network language modeling for code switching conversational speech. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8411–8415. IEEE.
- Heike Adel, Ngoc Thang Vu, and Tanja Schultz. 2013b. Combination of recurrent neural networks and factored language models for code-switching language modeling. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 206–211.
- Ashutosh Baheti, Sunayana Sitaram, Monojit Choudhury, and Kalika Bali. 2017. Curriculum design for code-switching: Experiments with language identification and language modeling with deep neural networks. *Proceedings of ICON*, pages 65–74.
- Hedi M Belazi, Edward J Rubin, and Almeida Jacqueline Toribio. 1994. Code switching and x-bar theory: The functional head constraint. *Linguistic inquiry*, pages 221–237.
- Susan Berk-Seligson. 1986. Linguistic constraints on intrasentential code-switching: A study of spanish/hebrew bilingualism. *Language in society*, 15(3):313–348.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of ibm model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648. Association for Computational Linguistics.
- Björn Gambäck and Amitava Das. 2014. On measuring the complexity of code-mixing. In *Proceedings of the 11th International Conference on Natural Language Processing, Goa, India*, pages 1–7. Citeseer.
- Saurabh Garg, Tanmay Parekh, and Preethi Jyothi. 2018. Code-switched language models using dual rnns and same-source pretraining. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3078–3083.
- Ying Li and Pascale Fung. 2012. Code-switch language model with inversion constraints for mixed language speech recognition. *Proceedings of COLING 2012*, pages 1671–1680.
- Zhaojiang Lin, Genta Indra Winata, and Pascale Fung. 2019. Learning comment generation by leveraging user-generated data. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7225–7229. IEEE.
- Yi Liu, Pascale Fung, Yongsheng Yang, Christopher Cieri, Shudong Huang, and David Graff. 2006. Hkust/mts: A very large scale mandarin telephone speech corpus. In *Chinese Spoken Language Processing*, pages 724–735. Springer.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Carol Myers-Scotton. 2001. The matrix language frame model: Development and responses. *Trends in Linguistics Studies and Monographs*, 126:23–58.
- Universiti Sains Malaysia Nanyang Technological University. 2015. Mandarin-english code-switching in south-east asia ldc2015s04. web download. philadelphia: Linguistic data consortium.
- Carol W Pfaff. 1979. Constraints on language mixing: intrasentential code-switching and borrowing in spanish/english. *Language*, pages 291–318.

- Shana Poplack. 1978. *Syntactic structure and social function of code-switching*, volume 2. Centro de Estudios Puertorriqueños, [City University of New York].
- Shana Poplack. 1980. Sometimes i'll start a sentence in spanish y termino en espanol: toward a typology of code-switching. *Linguistics*, 18(7-8):581–618.
- Shana Poplack. 2013. “sometimes i'll start a sentence in spanish y termino en español”: Toward a typology of code-switching. *Linguistics*, 51(Jubilee):11–14.
- Adithya Pratapa, Gayatri Bhat, Monojit Choudhury, Sunayana Sitaram, Sandipan Dandapat, and Kalika Bali. 2018. Language modeling for code-mixing: The role of linguistic theory based synthetic data. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1543–1553.
- Ofir Press and Lior Wolf. 2017. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 157–163.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083. Association for Computational Linguistics.
- Genta Indra Winata, Zhaojiang Lin, and Pascale Fung. 2019. Learning multilingual meta-embeddings for code-switching named entity recognition. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 181–186.
- Genta Indra Winata, Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018a. [Code-switching language modeling using syntax-aware multi-task learning](#). In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pages 62–67. Association for Computational Linguistics.
- Genta Indra Winata, Chien-Sheng Wu, Andrea Madotto, and Pascale Fung. 2018b. Bilingual character representation for efficiently addressing out-of-vocabulary words in code-switching named entity recognition. In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pages 110–114.
- LI Ying and Pascale Fung. 2014. Language modeling with functional head constraint for code switching speech recognition. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 907–916.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Shiyu Zhou, Linhao Dong, Shuang Xu, and Bo Xu. 2018. Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin chinese. In *Interspeech*.