

Meta-Transfer Learning for Code-Switched Speech Recognition

Genta Indra Winata^{†*}, Samuel Cahyawijaya*, Zhaojiang Lin,
Zihan Liu, Peng Xu, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)
Department of Electronic and Computer Engineering
The Hong Kong University of Science and Technology
[†]giwinata@connect.ust.hk

Abstract

An increasing number of people in the world today speak a mixed-language as a result of being multilingual. However, building a speech recognition system for code-switching remains difficult due to the availability of limited resources and the expense and significant effort required to collect mixed-language data. We therefore propose a new learning method, *meta-transfer learning*, to transfer learn on a code-switched speech recognition system in a low-resource setting by judiciously extracting information from high-resource monolingual datasets. Our model learns to recognize individual languages, and transfer them so as to better recognize mixed-language speech by conditioning the optimization on the code-switching data. Based on experimental results, our model outperforms existing baselines on speech recognition and language modeling tasks, and is faster to converge.

1 Introduction

In bilingual or multilingual communities, speakers can easily switch between different languages within a conversation (Wang et al., 2009). People who know how to code-switch will mix languages in response to social factors as a way of communicating in a multicultural society. Generally, code-switching speakers switch languages by taking words or phrases from the embedded language to the matrix language. This can occur within a sentence, which is known as *intra-sentential code-switching* or between two matrix language sentences, which is called *inter-sentential code-switching* (Heredia and Altarriba, 2001).

Learning a code-switching automatic speech recognition (ASR) model has been a challenging task for decades due to data scarcity and difficulty in capturing similar phonemes in different




Figure 1: Illustration of (a) joint training and (b) meta-transfer learning. The solid lines show the optimization path. The orange circles represent the monolingual source language, and the white circles represent the code-switching target language. The lower black circle in (b) is closer to T_{cs} than that in (a).

languages. Several approaches have focused on generating synthetic speech data from monolingual resources (Nakayama et al., 2018; Winata et al., 2019). However, these methods are not guaranteed to generate natural code-switching speech or text. Another line of work explores the feasibility of leveraging large monolingual speech data in the pre-training and applying fine-tuning on the model using a limited source of code-switching data, which has been found useful to improve the performance (Li et al., 2011; Winata et al., 2019). However, the transferability of these pretraining approaches is not optimized on extracting useful knowledge from each individual languages in the context of code-switching, and even after the fine-tuning step, the model forgets about the previously learned monolingual tasks.

In this paper, we introduce a new method, *meta-transfer learning*¹, to learn to transfer knowledge from source monolingual resources to a code-switching model. Our approach extends the model-

* These two authors contributed equally.

¹The code is available at <https://github.com/audioku/meta-transfer-learning>

agnostic meta learning (MAML) (Finn et al., 2017) to not only train with monolingual source language resources but also optimize the update on the code-switching data. This allows the model to leverage monolingual resources that are optimized to detect code-switching speech. Figure 1 illustrates the optimization flow of the model. Different from joint training, meta-transfer learning computes the first-order optimization using the gradients from monolingual resources constrained to the code-switching validation set. Thus, instead of learning one model that is able to generalize to all tasks, we focus on judiciously extracting useful information from the monolingual resources.

The main contribution is to propose a novel method to transfer learn information efficiently from monolingual resources to the code-switched speech recognition system. We show the effectiveness of our approach in terms of error rate, and that our approach is also faster to converge. We also show that our approach is also applicable to other natural language tasks, such as code-switching language modeling tasks.

2 Related Work

Meta-learning Our idea of learning knowledge transfer from source monolingual resources to a code-switching model comes from MAML (Finn et al., 2017). Probabilistic MAML (Finn et al., 2018) is an extension of MAML, which has better classification coverage. Meta-learning has been applied to natural language and speech processing (Hospedales et al., 2020). Madotto et al. (2019) extends MAML to the personalized text generation domain and successfully produces more person-consistent dialogue. Gu et al. (2018) and Qian and Yu (2019) and Lin et al. (2019) propose to apply meta-learning on low-resource learning. Yu et al. (2020) applies MAML to hypernym detection. Several applications have been proposed in speech applications, such as cross-lingual speech recognition (Hsu et al., 2019), speaker adaptation (Klejch et al., 2018, 2019), and cross-accent speech recognition (Winata et al., 2020).

Code-Switching ASR Li and Fung (2012) introduces a statistical method to incorporate a linguistic theory into a code-switching speech recognition system, and Adel et al. (2013a,b) explore syntactic and semantic features on recurrent neural networks (RNNs). Baheti et al. (2017) adapts effective curriculum learning by training a network

Algorithm 1 Meta-Transfer Learning

Require: $\mathcal{D}_{src}, \mathcal{D}_{tgt}$

Require: α, β : step size hyperparameters

- 1: Randomly initialize θ
 - 2: **while** not done **do**
 - 3: Sample batch data $\mathcal{D}^{tra} \sim (\mathcal{D}_{src}, \mathcal{D}_{tgt})$,
 $\mathcal{D}^{val} \sim \mathcal{D}_{tgt}$
 - 4: **for all** $\mathcal{D}_{\mathcal{T}_i}^{tra} \in \mathcal{D}^{tra}$ **do**
 - 5: Evaluate $\nabla_{\theta} \mathcal{L}_{\mathcal{D}_{\mathcal{T}_i}^{tra}}(f_{\theta})$ using $\mathcal{D}_{\mathcal{T}_i}^{tra}$
 - 6: Compute adapted parameters with gradient descent:
 $\theta'_{\mathcal{T}_i} = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{D}_{\mathcal{T}_i}^{tra}}(f_{\theta})$
 - 7: **end for**
 - 8: $\theta \leftarrow \theta - \beta \sum_i \nabla_{\theta} \mathcal{L}_{\mathcal{D}^{val}}(f_{\theta'_{\mathcal{T}_i}})$
 - 9: **end while**
-

with monolingual corpora of two languages, and subsequently training on code-switched data. Pratapa et al. (2018) and Lee et al. (2019) propose to use methods to generate artificial code-switching data using a linguistic constraint. Winata et al. (2018) proposes to leverage syntactic information to improve the identification of the location of code-switching points, and improve the language model performance. Finally Garg et al. (2018) and Winata et al. (2019) propose new neural-based methods using SeqGAN and pointer-generator (Pointer-Gen) to generate diverse synthetic code-switching sentences that are sampled from the real code-switching data distribution.

3 Meta-Transfer Learning

We aim to effectively transfer knowledge from source domains to a specific target domain. We denote our model by f_{θ} with parameters θ . Our model accepts a set of speech inputs $X = \{x_1, \dots, x_n\}$ and generates a set of utterances $Y = \{y_1, \dots, y_m\}$. The training involves a set of speech datasets in which each dataset is treated as a task \mathcal{T}_i . Each task is distinguished as either a source \mathcal{D}_{src} or target task \mathcal{D}_{tgt} . For each training iteration, we randomly sample a set of data as training \mathcal{D}^{tra} , and a set of data as validation \mathcal{D}^{val} . In this section, we present and formalize the method.

3.1 Setup

To facilitate the model to achieve a good generalization on the code-switching data, we sample the source dataset \mathcal{D}_{src} from monolingual English (*en*) and Chinese (*zh*) and code-switching (*cs*) corpora,

and choose the target dataset \mathcal{D}_{tgt} only from the code-switching corpus. The code-switching data samples between \mathcal{D}_{src} and \mathcal{D}_{tgt} are disjoint. In this case, we exploit the meta-learning update using meta-transfer learning to acquire knowledge from the monolingual English and Chinese corpora, and optimize the learning process on the code-switching data. Then, we slowly fine-tune the trained model to become closer to the code-switching domain by avoiding aggressive updates that can push the model to a worse position.

3.2 Meta-Transfer Learning Algorithm

Our approach extends the meta-learning paradigm to adapt knowledge learned from source domains to a specific target domain. This approach captures useful information from multiple resources to the target domain, and updates the model accordingly. Figure 1 presents the general idea of meta-transfer learning. The goal of the meta-transfer learning is not to focus on generalizing to all tasks, but to focus on acquiring crucial knowledge to transfer from monolingual resources to the code-switching domain. As shown in Algorithm 1, for each adaptation step on \mathcal{T}_i , we compute updated parameters $\theta'_{\mathcal{T}_i}$ via stochastic gradient descent (SGD) as follows:

$$\theta'_{\mathcal{T}_i} = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{D}_{\mathcal{T}_i}^{tra}}(f_{\theta}), \quad (1)$$

where α is a learning hyper-parameter of the inner optimization. Then, a cross-entropy loss $\mathcal{L}_{\mathcal{D}^{val}}$ is calculated from a learned model upon the generated text given the audio inputs on the target domain:

$$\mathcal{L}_{\mathcal{D}^{val}} = - \sum_{\mathcal{D}^{val} \sim \mathcal{D}_{tgt}} \log p(y_t | x_t; \theta'_{\mathcal{T}_i}). \quad (2)$$

We define the objective as follows:

$$\min_{\theta} \sum_{\mathcal{D}_{\mathcal{T}_i}^{tra}, \mathcal{D}^{val}} \mathcal{L}_{\mathcal{D}^{val}}(f_{\theta'_{\mathcal{T}_i}}) = \quad (3)$$

$$\sum_{\mathcal{D}_{\mathcal{T}_i}^{tra}, \mathcal{D}^{val}} \mathcal{L}_{\mathcal{D}^{val}}(f_{\theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{D}_{\mathcal{T}_i}^{tra}}(f_{\theta})}), \quad (4)$$

where $\mathcal{D}_{\mathcal{T}_i}^{tra} \sim (\mathcal{D}_{src}, \mathcal{D}_{tgt})$ and $\mathcal{D}^{val} \sim \mathcal{D}_{tgt}$. We minimize the loss of the $f_{\theta'_{\mathcal{T}_i}}$ upon \mathcal{D}^{val} . Then, we apply gradient descent on the meta-model parameter θ with a β meta-learning rate.

4 Code-Switched Speech Recognition

4.1 Model Description

We build our speech recognition model on a transformer-based encoder-decoder (Dong et al.,

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

Table 1: Data statistics of SEAME Phase II. CMI and SPF represents *code mixing index* and *switch-point fraction*, respectively.

2018; Winata et al., 2019). The encoder employs VGG (Simonyan and Zisserman, 2015) to learn a language-agnostic audio representation and generate input embeddings. The decoder receives the encoder outputs and applies multi-head attention to the decoder input. We apply a mask into the decoder attention layer to avoid any information flow from future tokens. During the training process, we optimize the next character prediction by shifting the transcription by one. Then, we generate the prediction by maximizing the log probability of the sub-sequence using beam search.

4.2 Language Model Rescoring

To further improve the prediction, we incorporate *Pointer-Gen LM* (Winata et al., 2019) in a beam search process to select the best sub-sequence scored using the softmax probability of the characters. We define $P(Y)$ as the probability of the predicted sentence. We add the pointer-gen language model $p_{lm}(Y)$ to rescore the predictions. We also include word count $wc(Y)$ to avoid generating very short sentences. $P(Y)$ is calculated as follows:

$$P(Y) = \alpha P(Y|X) + \beta p_{lm}(Y) + \gamma \sqrt{wc(Y)}, \quad (5)$$

where α is the parameter to control the decoding probability, β is the parameter to control the language model probability, and γ is the parameter to control the effect of the word count.

5 Experiments and Results

5.1 Dataset

We use SEAME Phase II, a conversational English-Mandarin Chinese code-switching speech corpus that consists of spontaneously spoken interviews and conversations (Nanyang Technological University, 2015). The data statistics and code-switching metrics, such as code mixing index (CMI) (Gambäck and Das, 2014) and switch-point

Model	CER
Winata et al. (2019)	32.76%
+ Pointer-Gen LM	31.07%
Only CS	34.51%
Joint Training (EN + ZH)	98.29%
+ Fine-tuning	31.22%
Joint Training (EN + CS)	34.77%
Joint Training (ZH + CS)	33.93%
Joint Training (EN + ZH + CS)	32.87%
+ Fine-tuning	31.90%
+ Pointer-Gen LM	31.74%
Meta-Transfer Learning (EN + CS)	32.35%
Meta-Transfer Learning (ZH + CS)	31.57%
Meta-Transfer Learning (EN + ZH + CS)	30.30%
+ Fine-tuning	29.99%
+ Pointer-Gen LM	29.30%

Table 2: Results of the evaluation in CER, a lower CER is better. *Meta-Transfer Learning* is more effective in transferring information from monolingual speech.

fraction (Pratapa et al., 2018) are depicted in Table 1. For monolingual speech datasets, we use HKUST (Liu et al., 2006) as the monolingual Chinese dataset, and Common Voice (Ardila et al., 2019) as the monolingual English dataset.² We use 16 kHz audio inputs and up-sample the HKUST data from 8 to 16 kHz.

5.2 Experiment Settings

Our transformer model consists of two encoder layers and four decoder layers with a hidden size of 512, an embedding size of 512, a key dimension of 64, and a value dimension of 64. The input of all the experiments uses spectrogram, computed with a 20 ms window and shifted every 10 ms. Our label set has 3765 characters and includes all of the English and Chinese characters from the corpora, spaces, and apostrophes. We optimize our model using Adam and start the training with a learning rate of $1e-4$. We fine-tune our model using SGD with a learning rate of $1e-5$, and apply an early stop on the validation set. We choose $\alpha = 1$, $\beta = 0.1$, and $\gamma = 0.1$. We draw the sample of the batch randomly with a uniform distribution every iteration.

We conduct experiments with the following approaches: (a) only CS, (b) joint training on EN + ZH, (c) joint training on EN + ZH + CS, and (d) meta-transfer learning. Then, we apply fine-tuning (b), (c), and (d) models on CS. We apply

²We downloaded the CommonVoice version 1 dataset from <https://voice.mozilla.org/>.

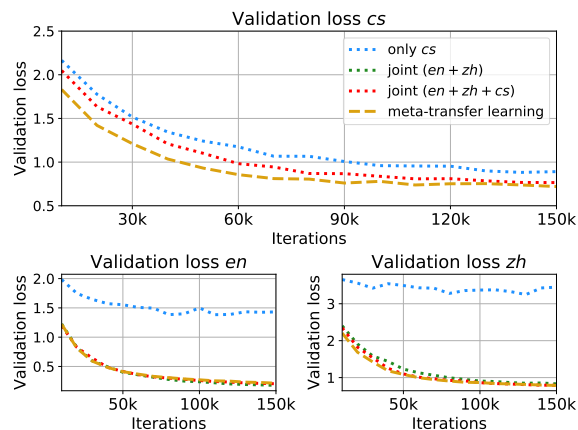


Figure 2: Validation loss per iteration. **Top:** validation loss on CS data, (joint (EN + ZH) is omitted because it is higher than the range), **bottom left:** validation loss on EN data, **bottom right:** validation loss on ZH data.

LM rescoring on our best model. We evaluate our model using beam search with a beam width of 5 and maximum sequence length of 300. The quality of our model is measured using character error rate (CER).

5.3 Results

The results are shown in Table 2. Generally, adding monolingual data EN and ZH as the training data is effective to reduce error rates. There is a significant margin between **only CS** and **joint training** (1.64%) or **meta-transfer learning** (4.21%). According to the experiment results, **meta-transfer learning** consistently outperforms the joint-training approaches. This shows the effectiveness of meta-transfer learning in language adaptation.

The fine-tuning approach helps to improve the performance of trained models, especially on the joint training (EN + ZH). We observe that joint training (EN + ZH) without fine-tuning cannot predict mixed-language speech, while joint training on EN + ZH + CS is able to recognize it. However, according to Table 3, adding a fine-tuning step badly affects the previous learned knowledge (e.g., EN: 11.84% \rightarrow 63.85%, ZH: 31.30% \rightarrow 78.07%). Interestingly, the model trained with meta-transfer learning does not suffer catastrophic forgetting even without focusing the loss objective to learn both monolingual languages. As expected, joint training on EN + ZH + CS achieves decent performance on all tasks, but it does not optimally improve CS.

The language model rescoring using Pointer-Gen LM improves the performance of the meta-transfer

Model	↑ Δ CS	↓ EN	↓ ZH
Only CS	-	66.71%	99.66%
Joint Training (EN + ZH)	-63.78%	11.84%	31.30%
+ Fine-tuning	3.29%	63.85%	78.07%
Joint Training (EN + ZH + CS)	1.64%	13.88%	30.46%
+ Fine-tuning	2.61%	57.56%	76.20%
Meta-Transfer Learning (EN + ZH + CS)	4.21%	16.22%	31.39%

Table 3: Performance on monolingual English CommonVoice test set (EN) and HKUST test set (ZH) in CER. Δ CS denotes the improvement on SEAME test set (CS) relative to the baseline model (Only CS).

Model	valid	test
Only CS [‡]	72.89	65.71
Joint Training (EN + ZH + CS)	70.99	63.73
+ Fine-tuning	69.66	62.73
Meta-Transfer Learning (EN + ZH + CS)	68.83	62.14
+ Fine-tuning	68.71	61.97

Table 4: Results on the language modeling task in perplexity. [‡] the results are from Winata et al. (2019).

learning model by choosing more precise code-switching sentences during beam search. Pointer-Gen LM improves the performance of the model, and outperforms the model trained only in CS by 5.21% and previous state-of-the-art by 1.77%.

Convergence Rate Figure 2 depicts the dynamics of the validation loss per iteration on CS, EN, and ZH. As we can see from the figure, meta-transfer learning is able to converge faster than only CS and joint training, and results in the lowest validation loss. For the validation losses on EN and ZH, both joint training (EN + ZH + CS) and meta-transfer learning achieve a similar loss in the same iteration, while only CS achieves a much higher validation loss. This shows that meta-transfer learning is not only optimized on the code-switching domain, but it also preserves the generalization ability to monolingual domains, as depicted in Table 3.

5.4 Language Modeling Task

We further evaluate our meta-transfer learning approach on a language model task. We simply take the transcription of the same datasets and build a 2-layer LSTM-based language model following the model configuration in Winata et al. (2019). To further improve the performance, we apply fine-tuning with an SGD optimizer by using a learning rate of 1.0, and decay the learning rate by 0.25x for every epoch without any improvement on the validation performance. To prevent the model from over-fitting, we add an early stop of 5 epochs.

As shown in Table 4, the meta-transfer learning approach outperforms the joint-training approach. We find a similar trend for the language model task results to the speech recognition task where meta-transfer learning without additional fine-tuning performs better than joint training with fine-tuning. Compared to our baseline model (Only CS), meta-transfer learning is able to reduce the test set perplexity by 3.57 points (65.71 \rightarrow 62.14), and the post fine-tuning step reduces the test set perplexity even further, from 62.14 to 61.97.

6 Conclusion

We propose a novel method, *meta-transfer learning*, to transfer learn on a code-switched speech recognition system in a low-resource setting by judiciously extracting information from high-resource monolingual datasets. Our model recognizes individual languages and transfers them so as to better recognize mixed-language speech by conditioning the optimization objective to the code-switching domain. Based on experimental results, our training strategy outperforms joint training even without adding a fine-tuning step, and it requires less iterations to converge.

In this paper, we have shown that our approach can be effectively applied to both speech processing and language modeling tasks. Finally, we will explore further the generability of our meta-transfer learning approach to more downstream multilingual tasks in our future work.

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.

References

- 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.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. 2019. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.
- 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.
- Linhao Dong, Shuang Xu, and Bo Xu. 2018. Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5884–5888. IEEE.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org.
- Chelsea Finn, Kelvin Xu, and Sergey Levine. 2018. **Probabilistic model-agnostic meta-learning**. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, pages 9537–9548, USA. Curran Associates Inc.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Yun Chen, Victor OK Li, and Kyunghyun Cho. 2018. Meta-learning for low-resource neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631.
- Roberto R Heredia and Jeanette Altabriba. 2001. Bilingual language mixing: Why do bilinguals code-switch? *Current Directions in Psychological Science*, 10(5):164–168.
- Timothy Hospedales, Antreas Antoniou, Paul Mi-caelli, and Amos Storkey. 2020. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*.
- Jui-Yang Hsu, Yuan-Jui Chen, and Hung-yi Lee. 2019. Meta learning for end-to-end low-resource speech recognition. *arXiv preprint arXiv:1910.12094*.
- Ondřej Klejch, Joachim Fainberg, and Peter Bell. 2018. Learning to adapt: A meta-learning approach for speaker adaptation. *Proc. Interspeech 2018*, pages 867–871.
- Ondřej Klejch, Joachim Fainberg, Peter Bell, and Steve Renals. 2019. Speaker adaptive training using model agnostic meta-learning. *arXiv preprint arXiv:1910.10605*.
- Grandee Lee, Xianghu Yue, and Haizhou Li. 2019. Linguistically motivated parallel data augmentation for code-switch language modeling. In *INTERSPEECH 2019*.
- 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.
- Ying Li, Pascale Fung, Ping Xu, and Yi Liu. 2011. Asymmetric acoustic modeling of mixed language speech. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5004–5007. IEEE.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Zihan Liu, Yan Xu, Cong Gao, and Pascale Fung. 2019. Learning to learn sales prediction with social media sentiment. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing*, pages 47–53.
- 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.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. **Personalizing dialogue agents via meta-learning**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5459, Florence, Italy. Association for Computational Linguistics.
- Sahoko Nakayama, Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2018. Speech chain for semi-supervised learning of japanese-english code-switching asr and tts. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 182–189. IEEE.

- Universiti Sains Malaysia Nanyang Technological University. 2015. Mandarin-english code-switching in south-east asia ldc2015s04. web download. philadelphia: Linguistic data consortium.
- 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.
- Kun Qian and Zhou Yu. 2019. [Domain adaptive dialog generation via meta learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2639–2649, Florence, Italy. Association for Computational Linguistics.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *ICLR*.
- Yapeng Wang, Patricia K Kuhl, Chunhui Chen, and Qi Dong. 2009. Sustained and transient language control in the bilingual brain. *NeuroImage*, 47(1):414–422.
- Genta Indra Winata, Samuel Cahyawijaya, Zihan Liu, Zhaojiang Lin, Andrea Madotto, Peng Xu, and Pascale Fung. 2020. Learning fast adaptation on cross-accented speech recognition. *arXiv preprint arXiv:2003.01901*.
- Genta Indra Winata, Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. [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, Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2019. Code-switched language models using neural based synthetic data from parallel sentences. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 271–280.
- Changlong Yu, Jialong Han, Haisong Zhang, and Wilfred Ng. 2020. Hypernymy detection for low-resource languages via meta learning. In *Proceedings of ACL*.