

INTapt: Information-Theoretic Adversarial Prompt Tuning for Enhanced Non-Native Speech Recognition

Eunseop Yoon^{1*}, Hee Suk Yoon^{1*}, John Harvill²,
Mark Hasegawa-Johnson², and Chang D. Yoo^{1†}

¹Korea Advanced Institute of Science and Technology (KAIST)

²University of Illinois at Urbana-Champaign (UIUC)

{esyoon97, hskyoony, cd_yoo}@kaist.ac.kr

{harvill2, jhasegaw}@illinois.edu

Abstract

Automatic Speech Recognition (ASR) systems have attained unprecedented performance with large speech models pre-trained based on self-supervised speech representation learning. However, these pre-trained speech models suffer from representational bias as they tend to better represent those prominent accents (i.e., native (L1) English accent) in the pre-training speech corpus than less represented accents, resulting in a deteriorated performance for non-native (L2) English accents. Although there have been some approaches to mitigate this issue, all of these methods require updating the pre-trained model weights. In this paper, we propose Information Theoretic Adversarial Prompt Tuning (INTapt), which introduces prompts concatenated to the original input that can re-modulate the attention of the pre-trained model such that the corresponding input resembles a native (L1) English speech without updating the backbone weights. INTapt is trained simultaneously in the following two manners: (1) adversarial training to reduce accent feature dependence between the original input and the prompt-concatenated input and (2) training to minimize CTC loss for improving ASR performance to a prompt-concatenated input. Experimental results show that INTapt improves the performance of L2 English and increases feature similarity between L2 and L1 accents.

1 Introduction

Self-supervised learning has improved input data representation without requiring extensive human-labeled data (He et al., 2019; Zhang et al., 2022). Based on this advancement, powerful pre-trained models providing high-performing representations for various data types (e.g., text, images, and audio) have been proposed. For instance, in speech, self-supervised pre-trained models such as HuBERT (Hsu et al., 2021) have advanced state-of-the-

*Equal contribution

†Corresponding author

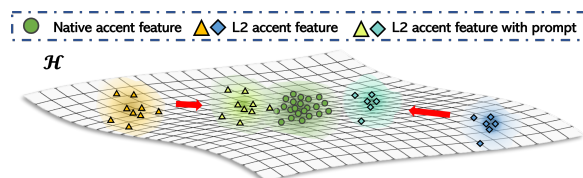


Figure 1: Illustration of a hypothetical accent feature space (\mathcal{H}). Distinctive accent features between native and L2 accents lead to degraded performance of ASR systems on L2 accents. INTapt concatenates a prompt to the input space to reduce this distinction.

art performance of automatic speech recognition (ASR).

However, one major challenge in using pre-trained speech models for ASR is the representational bias towards prominent accents present in the dataset during pre-training. Consequently, there will be a disparity in ASR performance between native and non-native speakers. More specifically, pre-training using a large dataset such as the LibriSpeech (Panayotov et al., 2015), which comprises a large proportion of utterances from native (L1) English speakers, leads to a less satisfactory recognition rate for non-native (L2) English accented speech. This phenomenon can curtail the effectiveness of current high-performing ASR systems for real-world applications.

There have been several ways to address this issue, including fine-tuning the model on diverse accents (Winata et al., 2019; Shibano et al., 2021), having a separate model for each accent (Yang et al., 2018) or using regularization losses that guide the fine-tuning process to achieve robustness to accents (Chen et al., 2020; Gao et al., 2022), all of which require updating the pre-trained model.

We propose a different solution for improving L2 speech recognition in transformer-based speech models that introduces a small number of learnable parameters into the input space while keeping the backbone weights of the model untouched.

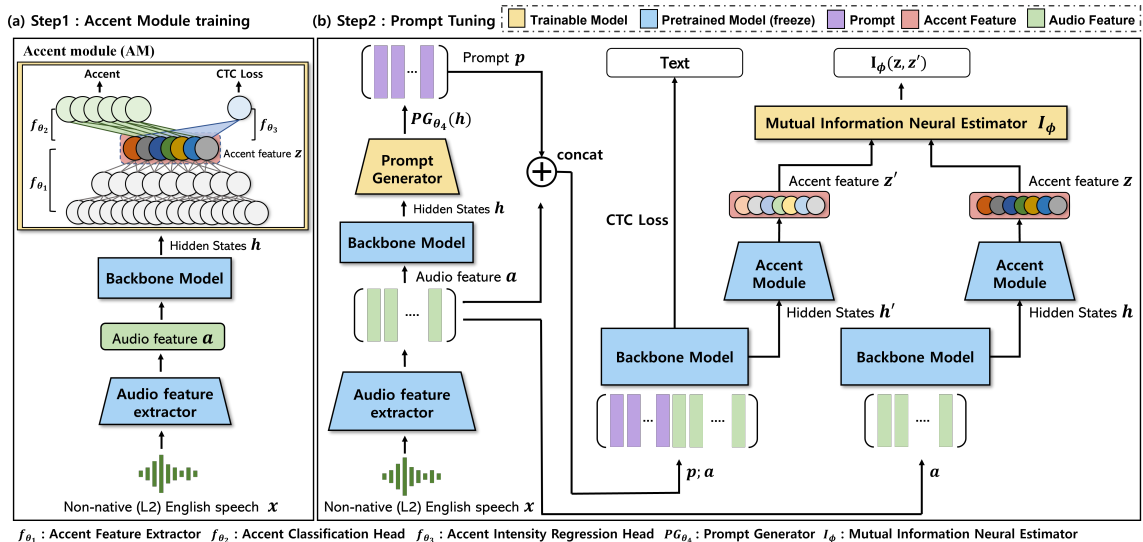


Figure 2: Overview of INTapt. INTapt incorporates a two-step training process where the first step involves training the Accent Module to get the accent feature of a particular input speech and the second step involves training the Prompt Generator capable of making the non-native (L2) English speech input have a better ASR performance by re-modulating the attention of the Backbone Model so that it resembles the accent of a native (L1) English speech.

Our approach is guided by Information-Theoretic Adversarial Learning; thus, we refer to it as INTapt (Information-Theoretic Adversarial Prompt Tuning). INTapt aims to introduce auxiliary embeddings (i.e., prompt) concatenated to the original input, which can re-modulate the attention and adapt the pre-trained weights so that the corresponding input looks like speech with an accent seen during pre-training (Figure 1). To achieve this, INTapt incorporates (1) adversarial training, which tries to minimize the mutual information between the accent feature of the original input and that obtained by concatenating the prompt embeddings in front of the initial input, and (2) CTC loss training to improve the ASR performance of the prompt-concatenated input. Essentially the prompt is trained such that the accent of the concatenation is pushed away from the input accent and the concatenation achieves native CTC loss performance. Unlike the previous use-case of prompts in NLP or Computer vision (CV), where a single prompt embedding is learned for each discrete task or input domain, the intensity of an accent is continuous. Thus, we propose an input-dependent prompt embedding by training a prompt generator that outputs an input-specific prompt. Through extensive experiments, we show that the proposed dual objectives of INTapt not only lead to better performance on L2 English accents but result in a higher similarity between the accent feature of the prompt-concatenated input and that of L1 English accents.

2 INTapt

Figure 2 depicts the overall process of INTapt. INTapt incorporates a two-step training process. In the first step, we train an Accent Module (AM) capable of isolating the accent feature from a given audio feature a of an input speech x . In the second step, we train a Prompt Generator (PG), which outputs a prompt p for a given audio feature a , using two objectives: (1) Minimize the mutual information between the accent feature z' and z , where the former is obtained using the prompt-concatenated input ($p; a$) and the latter is obtained from the original audio feature a , (2) Minimize CTC loss to improve the ASR performance of the input ($p; a$).

2.1 Accent Module (AM)

Since our method requires direct access to the isolated accent feature of the corresponding audio feature input, we propose an Accent Module (AM) capable of extracting the accent feature z from the input a . The module consists of an accent feature extractor f_{θ_1} which is trained with an accent classification head f_{θ_2} to isolate the accent feature and an accent intensity regression head f_{θ_3} to capture the intensity of the accent into the obtained feature.

Accent Classification Head The role of the accent classification head f_{θ_2} is to isolate the accent feature of a given speech¹. Given the hidden state

¹We show in Appendix B that the proposed way effectively isolate the accent feature from other features.

representation \mathbf{h} of an audio feature input \mathbf{a} , the feature extractor outputs the accent feature (i.e., $\mathbf{z} = f_{\theta_1}(\mathbf{h})$) and the accent classification head f_{θ_2} tries to assign it to the correct accent label y .

Accent Intensity Regression Head The intensity of an accent could vary among different people even though there are in the same L2 group, and it could also vary between utterances from the same speaker. Thus, an accent intensity regression head is introduced to incorporate the accent intensity into the obtained accent feature \mathbf{z} . Based on the assumption that the intensity of the accent affects ASR performance, making the accent intensity regression head predict the CTC loss², obtained by inputting the corresponding speech into the backbone speech model, will allow the extracted accent feature \mathbf{z} to capture the intensity of the accent.

Given a batch B , the training of the Accent Module with the two aforementioned heads could be summarized as:

$$\min_{\theta_1, \theta_2} \frac{1}{|B|} \sum_{i \in B} -\log p(y_i | f_{\theta_2}(f_{\theta_1}(h_i))) + \lambda \min_{\theta_1, \theta_3} \frac{1}{|B|} \sum_{i \in B} [f_{\theta_3}(f_{\theta_1}(h_i)) - \text{CTC}(x_i)]^2 \quad (1)$$

2.2 Prompt Generator (PG)

Building on the success of prompts in NLP (Liu et al., 2021; Li and Liang, 2021) and CV (Dosovitskiy et al.), we introduce a prompt tuning method to improve the ASR performance for L2 English speech by efficiently utilizing a pre-trained model that already shows good performance for L1 English speech. In contrast to traditional NLP or CV applications, where a single, discrete prompt embedding is learned for each specific task or input domain, the intensity of an accent is continuous. To address this, we propose an input-dependent prompt embedding by training prompt generator PG_{θ_4} that generates an input-specific prompt guided by Information-Theoretic Adversarial Learning. More specifically, given a hidden state $\mathbf{h} = [h_1, h_2, \dots, h_L]$ with length L we produce a prompt of length L' ,

$$\mathbf{p} = PG_{\theta_4}(\mathbf{h}) \quad (2)$$

Mutual Information Minimization Mutual Information measures the co-dependence between

²Connectionist Temporal Classification (CTC) (Graves et al., 2006) is the primary loss used to train deep neural networks in speech recognition.

two random variables X and Y . Belghazi et al. (2018) recently proposed a gradient descent based method for estimating this property, allowing the use of neural networks for the estimation of mutual information between high dimensional random variables. The estimation is done using a neural network parameterized by ϕ as below:

$$I(X, Y) \geq I_{\phi}(X, Y), \quad (3)$$

where maximizing $I_{\phi}(X, Y)$ provides a tight lower bound of the original mutual information $I(X, Y)$. We use this to adversarially train the prompt generator PG_{θ_4} to minimize the mutual information between the accent feature of the original L2 speech input and the prompt-concatenated input.

CTC Loss Minimization We train the prompt generator PG_{θ_4} to minimize the CTC loss obtained for the prompt-concatenated input $(\mathbf{p}; \mathbf{a})$.

The two minimization objectives wrt. the prompt generator, along with the maximization objective wrt. the Mutual Information Neural Estimator, are done jointly in the second training step (Equation 4). We show in Section 3.2 and 4 that the aforementioned objectives not only improve the ASR performance of L2 speech but also effectively make it resemble the accent feature of the L1 speech.

$$\min_{\theta_4} \max_{\phi} \frac{1}{|B|} \sum_{i \in B} \text{CTC}(\mathbf{p}_{\theta_4}; \mathbf{a}) + \lambda I_{\phi}(\mathbf{z}'_{\theta_4}, \mathbf{z}) \quad (4)$$

3 Experiments

3.1 Experimental setting

Dataset We use the L2-ARCTIC (Zhao et al., 2018), which is a speech corpus of non-native (L2) English speakers - Mandarin (ZH), Hindi (HI), Vietnamese (VI), Korean (KO), Spanish (ES), and Arabic (AR). Each L2 group contains two male and two female speakers, and all the speakers read the same 1132 texts. The train/dev/test set is configured by dividing the data into 0.8/0.1/0.1 splits with no overlapping texts between each splits. Additionally, since we would like to simulate a natural data collection situation where the amount of data varies across groups (Yoon et al., 2022), we randomly divided the training data into More Frequent Accent (MFA) (ZH, HI), Less Frequent Accent (LFA) (VI, KO), and Unseen Accent (UA) (ES, AR) - For MFA we keep all the training data, for LFA we keep half of the data, and for UA we remove all the training data.

Backbone	#.params	Methods	MFA		LFA		UA		ALL
			ZH	HI	VI	KO	ES	AR	
HUBERT _{Large}	-	Backbone	18.71	8.80	25.8	10.98	14.12	14.92	15.55
	315M	+Finetune	15.46	7.91	22.26	9.95	14.19	13.94	13.95
	12.5M	+Prompt _{ctc}	13.93	7.20	21.93	9.69	12.64	12.38	12.96
	12.9M	+INTapt	13.09	6.64	21.25	8.97	12.18	11.92	12.34
HUBERT _{XLarge}	-	Backbone	17.03	7.48	26.02	10.49	13.65	13.52	14.69
	958M	+Finetune	15.49	7.53	24.09	10.02	13.48	12.56	13.86
	19.7M	+Prompt _{ctc}	13.02	7.31	19.26	8.05	10.46	10.38	11.41
	19.9M	+INTapt	11.67	6.63	18.41	7.17	10.44	10.55	11.00

Table 1: Comparison of WER (%) (lower is better) on the created subset of L2-ARCTIC (MFA, LFA, UA). #.params denote the number of parameters that were updated for training. The std. are reported in Table 3 of Appendix C.2.

Models For the backbone pre-trained speech models we try two different settings, HuBERT_{Large} and HuBERT_{XLarge} (Hsu et al., 2021). We consider three different training situations: 1) **Finetune** denotes a standard finetuning method where we update the pre-trained model weights to minimize the CTC loss, 2) **Prompt_{ctc}** is the case of training the prompt generator without the minimization of mutual information, and 3) **INTapt** trains the prompt generator with our proposed objective in equation 4. We include the training details in Appendix A.

3.2 Results

Table 1 shows the Word Error Rate (WER) across different L2 groups on the ASR task. We find that the performance improvement of the prompt tuning approaches (Prompt_{ctc} and INTapt) are more significant compared to standard finetuning despite updating small number of parameters (2-4%). INTapt shows the lowest WER on all L2 groups, obtaining 12.34% for HuBERT_{Large} and 11.00% for HuBERT_{XLarge} on the aggregated all speakers, outperforming the finetuned by 1.62%p and 2.86%p, respectively³. This conforms to the previous findings (Lester et al., 2021) that larger model size can benefit more from prompt tuning methods.

In Table 2, we report the WER on LibriSpeech (Panayotov et al., 2015) test-clean and test-other, which consists mainly of L1 speech. Compared with the backbone model, the WER after finetuning increased by 5.81%p. However, since Prompt_{ctc} and INTapt does not change the backbone weights, the WER on test-all increased only by 0.48%p and 0.37%p, respectively. This shows one of the key

³We show some examples of improved ASR results using INTapt in Appendix C.1

Methods	test-clean	test-other	test-all
Backbone	2.15	4.42	3.29
+Finetune	8.10	10.08	9.10
+Prompt _{ctc}	2.56	4.93	3.77
+INTapt	2.41	4.94	3.66

Table 2: WER (%) (lower is better) on LibriSpeech. test-all denotes the aggregation of test-clean and test-other.

benefits of prompt tuning methods in that it only slightly degrades the performance of the backbone model on tasks it already excels at while improving performance on others.

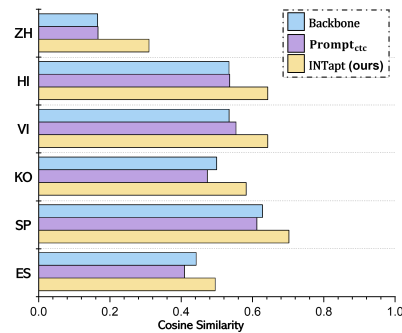


Figure 3: Cosine similarity between L1 accent feature and L2 accent features obtained from different methods.

4 Ablation Study

We analyze whether INTapt allows the L2 speech input to resemble the accent of L1 speech (Figure 3). Using the Accent Module, we extract the L1 accent feature, and L2 accent features obtained using the Backbone model, Prompt_{ctc}, and INTapt. INTapt showed the highest cosine similarity for all L2 groups, meaning that INTapt effectively adjusts the attention of the pre-trained model so that L2 speech resemble L1 speech in terms of accent.

5 Conclusion

We introduced Information Theoretic Adversarial Prompt Tuning (INTapt) for improving non-native ASR performance. To achieve this, INTapt re-modulates the attention of the pre-trained speech models by concatenating input-dependent prompt embeddings to the original input, without updating the model weights. Throughout the experiment, we show that INTapt is capable of outperforming standard finetuning of the pre-trained model on L2 speech, without degradation on L1 speech, by allowing the L2 input to resemble a L1 accent.

Limitations

INTapt adopts a prompt tuning method which utilizes the inherent information from pre-trained models that already shows good ASR performance on L1 English speakers. Therefore, in order to apply our method we need a pre-trained model that already has good performance on a specific task which might not be available for other languages. Also, our method might potentially need sufficiently large pre-trained model size in order for prompt to utilize the internal information of the model.

Ethics Statement

Since pre-trained speech models usually show better performance on native (L1) speech automatic speech recognition (ASR) due to the nature of the pre-training data used, this work have contributed to improve the ASR performance for non-native (L2) English speakers and mitigating the performance gap between them. This has the potential to construct a fair ASR machine well-operating not only on L1 English speakers but L2 speakers, which is an important feature to have for its deployment in real-life. Additionally, since we utilize the pre-trained model, it is possible to have ethical issue depending on the pre-trained model.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00951, Development of Uncertainty-Aware Agents Learning by Asking Questions), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea govern-

ment(MSIT) [NO.2021-0-02068, Artificial Intelligence Innovation Hub (Seoul National University)], and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-0-00184, Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics).

References

- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. 2018. [Mine: Mutual information neural estimation](#).
- Yi-Chen Chen, Zhaojun Yang, Ching-Feng Yeh, Mahaveer Jain, and Michael L Seltzer. 2020. Aipnet: Generative adversarial pre-training of accent-invariant networks for end-to-end speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6979–6983. IEEE.
- A Dosovitskiy, L Beyer, A Kolesnikov, D Weissenborn, et al. An image is worth 16x16 words: Transformers for image recognition at scale.
- Heting Gao, Xiaoxuan Wang, Sunghun Kang, Rusty John Lloyd Mina, Dias Issa, John Harvill, Leda Sari, Mark A. Hasegawa-Johnson, and Chang Dong Yoo. 2022. Seamless equal accuracy ratio for inclusive ctc speech recognition. *Speech Commun.*, 136:76–83.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. [Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks](#). In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, page 369–376, New York, NY, USA. Association for Computing Machinery.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2019. [Momentum contrast for unsupervised visual representation learning](#). *CoRR*, abs/1911.05722.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushan Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Toshiko Shibano, Xinyi Zhang, Mia Taige Li, Haejin Cho, Peter Sullivan, and Muhammad Abdul-Mageed. 2021. Speech technology for everyone: Automatic speech recognition for non-native english. In *Proceedings of The Fourth International Conference on Natural Language and Speech Processing (ICNLSP 2021)*, pages 11–20.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- 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 (RepLANLP-2019)*, pages 181–186, Florence, Italy. Association for Computational Linguistics.
- Xuesong Yang, Kartik Audhkhasi, Andrew Rosenberg, Samuel Thomas, Bhuvana Ramabhadran, and Mark Hasegawa-Johnson. 2018. Joint modeling of accents and acoustics for multi-accent speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Hee Suk Yoon, Eunseop Yoon, John Harvill, Sunjae Yoon, Mark Hasegawa-Johnson, and Chang Yoo. 2022. [SMSMix: Sense-maintained sentence mixup for word sense disambiguation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1493–1502, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chaoning Zhang, Kang Zhang, Trung Xuan Pham, Axi Niu, Zhinan Qiao, Chang Dong Yoo, and In-So Kweon. 2022. Dual temperature helps contrastive learning without many negative samples: Towards understanding and simplifying moco. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14421–14430.
- Guanlong Zhao, Sinem Sonsaat, Alif Silpachai, Ivana Lucic, Evgeny Chukharev-Hudilainen, John Levis, and Ricardo Gutierrez-Osuna. 2018. L2-arctic: A non-native english speech corpus. In *INTER-SPEECH*, pages 2783–2787.

A Experiment Details

Model Architecture We use a 3-layer Multi Layer Perceptron (MLP) for the Accent Feature Extractor and 1-layer, 3-layer MLP for the Accent Classification Head and the Accent Intensity Regression Head in the Accent Module, respectively. The Prompt Generator (PG) is composed of a single layer transformer. Since we adopt the transformer architecture for PG, its maximum output length is same as the length L of the input audio feature a . The specific length of prompt can be set by taking the first L' output embeddings from the front of the transformer output. For the Mutual Information Neural Estimator (MINE), we use a 3-layer MLP as well.

Training Details For pre-processing L2-ARCTIC, we utilized the huggingface resampling tool ⁴ to downsample the audio files from 44.1kHz to 16kHz. The hidden state representation obtained from the 3rd layer of the backbone model is used as the input to the Accent Module and Prompt Generator for both HuBERT_{large} and HuBERT_{XLarge} . The dimension of the accent feature a is set to $d = 256$, the length of prompt L' is 40, and the dimension of the prompt is set to 1024 and 1280, same as that of the input embedding size for HuBERT_{large} and HuBERT_{XLarge} , respectively. We use the AdamW optimizer (Loshchilov and Hutter, 2019) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 8$, and weight decay $\lambda = 0.005$ with different learning rates for all trainable model (i.e., AM, PG, MINE, finetuned backbone). The learning rate used for both AM and MINE is $1e-3$, and $5e-6$, $1e-4$, $1e-4$ are used for Fintune, prompt_{ctc} , and INTapt, respectively. For all the methods, the batch size is set to 16 for HuBERT_{large} and 8 for HuBERT_{XLarge} . We use $\lambda = 0.5$ for Equation 1 and $\lambda = 0.003$ for Equation 4. The best model is selected by the lowest WER on the validation set. All experiments was done on NVIDIA Quadro RTX 8000.

B Accent Feature Isolation

We visually analyze the accent feature extracted from AM to validate that the feature does not contain any other information except accent. We plot the 2-D representation of extracted accent feature using t-SNE (van der Maaten and Hinton, 2008) with the label for gender for three L2 groups (i.e., HI, KO, ES). Figure 5 shows that the scatter points are distinctive between L2 groups but difficult to distinguish gender, which means our AM successfully isolates the accent feature from audio.

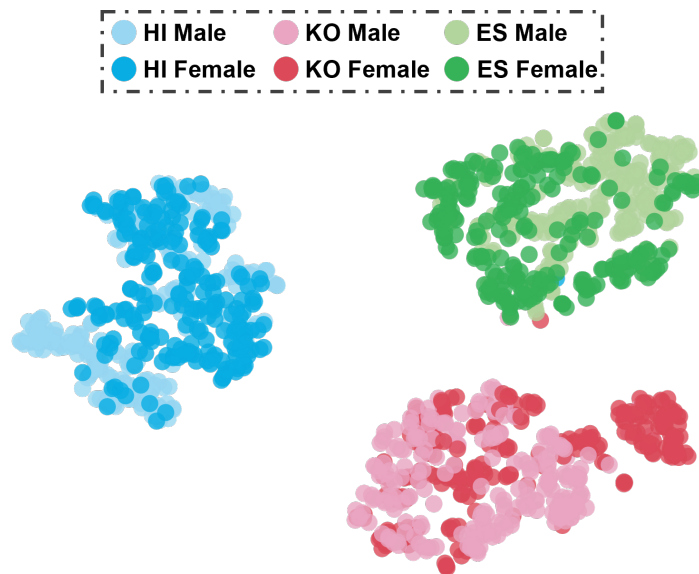


Figure 4: Latent space visualization showing that Accent Module extracts isolated accent feature

⁴https://huggingface.co/docs/datasets/audio_process

C Results of INTapt

C.1 Examples

In Figure 5, we show some examples of improved speech recognition results using INTapt. HuBERT_{Large} is used as backbone model for both cases and the INTapt having lowest WER on val set is selected to predict samples. **Red** represents wrong transcription output from the backbone pre-trained HuBERT_{Large}, and **green** represents enhanced recognition through the use of INTapt to re-modulate the backbone attention.

Ground Truth	Backbone Prediction	INTapt prediction (ours)
Then he shouted shut up	Then he shouldered shut up	Then he shouted shut up
You live on an income which your father earned	You live on an income which your father owned	You live on an income which your father earned
Straight out they swam their heads growing smaller and smaller	Straight out they swam their hats growing smaller and smaller	Straight out they swam their heads growing smaller and smaller
Seventeen no eighteen days ago	Seventeen no eighteen days acle	Seventeen no eighteen days ago
This is my fifth voyage	This is my fit for yet	This is my fifth voyage
But they make the mistake of ignoring their own duality	But they make the mistake of ignoring their unduality	But they make the mistake of ignoring their own duality
The mob came on but it could not advance	The mob came on but it could not at once	The mob came on but it could not advance
The big eyed clucking moose birds were most annoying	The big eyed clucking moose bords were most annoying	The big eyed clucking moose birds were most annoying
Philip dropped back into his chair	Philip jropped back into his chair	Philip dropped back into his chair
Would you be satisfied with that one hundredth part of me	Would you be satisfied wi dhat one hondref part o me	Would you be satisfied with that one hundredth part of me
He moved his position and the illusion was gone	Removed his position and the evolution was gone	He moved his position and the illusion was gone

Figure 5: Examples of enhanced recognition through the use of INTapt. **Red** represents the wrong transcription output prediction from the backbone model and **green** represents the correct output prediction using INTapt.

C.2 Standard Deviation of Results

Backbone	#.params	Methods	MFA		LFA		Unseen Accent		Avg.
			ZH	HI	VI	KO	ES	AR	
HUBERT _{Large}	315M	+Finetune	0.31	0.70	0.51	0.50	0.73	0.38	0.36
	12.5M	+Prompt _{ctc}	0.51	0.56	0.99	0.74	0.78	0.38	0.30
	12.9M	+INTapt	0.66	0.72	0.73	0.69	0.59	0.27	0.13
HUBERT _{XLarge}	958M	+Finetune	0.41	0.21	1.79	0.06	0.33	0.12	0.29
	19.7M	+Prompt _{ctc}	0.15	0.64	0.19	0.51	0.24	0.56	0.12
	19.9M	+INTapt	0.33	0.31	0.41	0.30	0.25	0.12	0.18

Table 3: Standard deviation values for the experimental results in Table 1. The values were obtained by running the same experiments with five different random seeds.

In Table 3, we report the standard deviation of the results in Table 1 with five different random seeds. As the backbone experiment in Table 1 is obtained without any training, we do not contain the standard deviation for those.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
The limitations of our work is contained in 'Limitation' section after conclusion.
- A2. Did you discuss any potential risks of your work?
It is discussed in 'Ethics statement' section.
- A3. Do the abstract and introduction summarize the paper's main claims?
The summarization of our main claims are included in both abstract and Introduction (sec1)
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

In Section 3.1 and 3.2

- B1. Did you cite the creators of artifacts you used?
In Section 3.1 and 3.2
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We provide all references for datasets and models, and the licenses or terms can be checked in the references.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
In section 3.1 and 3.2, we implicitly mention that datasets and models are used with thier own intended use.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We utilized publically-open dataset, and the contents are all from audio book. So we don't need to contain any step to anonymize contetnts. Also we don't introduce or generate new data samples and this is mentioned in Ethics statement section.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We clearly include the language that we use for speech recognition and the composition of L2 groups for non-native accented English speech samples.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
In section 3 dataset part

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

We include training details in Appendix A

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

We report the number of parameters for the proposed method in results, the GPU that we used for the experiments in appendix A.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

In appendix A, Experiment details part

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

In section 3.2, the results are avg. of 5 runs using different seeds, and we have std. in Appendix C3

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

In appendix A, we contains the link for the package that we used for data pre-processing

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.