



Black-box Adaptation of ASR for Accented Speech

Kartik Khandelwal, Preethi Jyothi, Abhijeet Awasthi, Sunita Sarawagi

Indian Institute of Technology Bombay, Mumbai, India

{kartikk,pjyothi,awasthi,sunita}@cse.iitb.ac.in

Abstract

We introduce the problem of adapting a black-box, cloud-based ASR system to speech from a target accent. While leading online ASR services obtain impressive performance on mainstream accents, they perform poorly on sub-populations — we observed that the word error rate (WER) achieved by Google’s ASR API on Indian accents is almost twice the WER on US accents. Existing adaptation methods either require access to model parameters or overlay an error correcting module on output transcripts. We highlight the need for correlating outputs with the original speech to fix accent errors. Accordingly, we propose a novel coupling of an open-source accent-tuned local model with the black-box service where the output from the service guides frame-level inference in the local model. Our fine-grained merging algorithm is better at fixing accent errors than existing word-level combination strategies. Experiments on Indian and Australian accents with three leading ASR models as service, show that we achieve upto 28% relative reduction in WER over both the local and service models.

Index Terms: Black box ASR systems, accented speech recognition, adaptation.

1. Introduction

The emergence of cloud-based AI services, for tasks like machine translation and speech recognition, have greatly increased the accessibility of machine learning. These services are powered by sophisticated engines and trained on large proprietary datasets. The internals of these engines are often not exposed to clients. Often a client’s input comes from a different domain than the training domain of the server. The existing fix of re-training to adapt to new target domains is not an option in this case. This leads us to our problem of *Black-box Adaptation*.

In this work, the task of interest is automatic speech recognition (ASR) in English, and the domains correspond to different accents. While leading online services like the Google ASR API [1] attain superior performance on high-resource English accents, they perform poorly on a large number of under-represented English accents. The API gives word error rates (WERs) of 23% or higher on datasets in Australian and Indian accents, as opposed to a WER of 13.2% on US accents. As ASR systems start getting deployed in several critical applications, it is increasingly imperative to design light-weight methods of accent adaptation to provide fair access to users of all regions and ethnicity [2]. Existing methods of adapting ASR models to a specific accent [3, 4, 5, 6, 7, 8], require modifying model parameters, which is not an option in the black-box setting.

One could implement black-box adaptation in the form of an error-correction model to alter the outputs from the service [9]. When the mismatch is in the language model, error correction models using a domain-specific language model have been proposed before [10]. However, for recovering from accent errors an error correction model would need to correlate the

service’s transcript with the original speech. To handle speech, the model would in turn need to incorporate an ASR system. This leads us to switching our perspective, so that black-box adaptation amounts to building a local ASR system which is retargeted to correct accent errors in the service’s output.

The local ASR system would be an open-sourced ASR architecture like DeepSpeech2 [11] pretrained on a publicly available corpus like the US-accented LibriSpeech [12] corpus but further finetuned using a small amount of data in the target accent. Typically the local model would be less accurate than the service in all parts except the parts with systematic accent differences. If the outputs from the local and service models are combined via standard combination approaches at the word or transcript-level [13, 14], we obtain only limited improvements in accuracy over the service. In other words, if the local system were also to be used as a black-box, we would not obtain the performance improvements we seek.

Hence, we exploit our white-box access to the local system. Our idea, at a high-level, is to use the transcript obtained from the service to *guide the inference* of the local ASR system. Our guided inference algorithm (named FineMerge) aligns the characters in the service with input frames using a Viterbi-like decoding and then selectively modifies the frame-level distribution of the local model. Our fine-grained merging step is easy to plug in existing speech pipelines, fast during inference, and specifically tailored to fixing accent errors — we often recover words that were absent from stand-alone outputs of both local and service models. Experiments on different service APIs on two different English accents show that FineMerge provides significant reduction in WER over either the local or service models, and existing methods of combining them at the word-level. To summarize, our overall contribution in this paper are:

1. We introduce the problem of black-box accent adaptation of ASR service APIs.
2. We propose an efficient coupling of a local white-box model with a black-box service to accent adapt with limited labeled data without incurring the cost of accessing the service during training.
3. We design a novel guided inference algorithm on the local model that is specifically tailored to correct focused accent errors in an otherwise strong service API.
4. We evaluate our algorithm on two accents and three service APIs. Our approach provides up to 28% reduction in relative WER over both local and service models. Existing methods based on rescoring N-best lists or combining outputs at the word-level are not as effective.

2. Related Work

Accent Adaptation in Speech. Accent adaptation in speech has been a problem of long standing interest. One category of methods attempt to create accent invariant systems and range from early approaches that simply merged data from multiple

accents for training a single model [15] to more recent work that uses adversarial learning objectives to extract accent-invariant feature representations from speech [16, 17]. A second category of methods are accent dependent methods that adapt to the speaker’s accent. Early approaches were HMM-based acoustic model adaptation and pronunciation model augmentation with accent-specific pronunciations [18, 19]. Within neural models, accent adaptation was achieved via accent-specific output layers [3, 4] and hierarchical models in a multitask learning setting [8]. A more recent work jointly learns an accent classifier and accent-dependent models [5, 6, 7]. Our method is also accent dependent but we need to adapt a black-box service model. We build local accent-adapted ASR systems, which are in turn guided by inference by service predictions.

Black box ASR Systems. Speech transcription services have seen widespread use in recent years. However, the underlying ASR systems in these services are black box systems. Adapting such models to a client’s needs would be of great utility but prior work in this area is sparse. [20] shows how to optimize black box ASR systems and [21] shows how to improve confidence estimates produced by such black-box systems. Another closely related work [10] is to use a domain-specific language model and a semantic parser to rescore the hypotheses from a black-box ASR system. Unlike their method, we achieve a more fine-grained integration of our client model with the service.

System Combination Approaches. Ours can be viewed as a type of system combination approach which has seen wide use in ASR. ROVER (Recognizer Output Voting Error Reduction) [13] is one of the most popular techniques that first combines predictions from different systems using an alignment step followed by a weighted voting step. Prior work on dialectal speech recognition [14] observed that using the best output from a dialect-specific model is more accurate than techniques like ROVER. Unlike ROVER that considers each individual system as a black-box, our method that leverages white-box access to a local accent-adapted ASR system is more targeted to correct accent errors and ultimately more accurate.

3. Our Approach

Given an audio signal \mathbf{x} , we invoke the service model \mathcal{S} on \mathbf{x} and get the transcript \mathbf{s} comprising of tokens s_1, \dots, s_k , along with token-level confidences p_1, \dots, p_k . In addition the client can invoke a local white-box model \mathcal{C} that has been trained/fine-tuned on limited accented labeled data. On input \mathbf{x} , let $\mathbf{c} = c_1, \dots, c_r$ denote the transcript from the local model \mathcal{C} with token-level confidences $\mathbf{q} = q_1, \dots, q_r$. In general the number of tokens in the two outputs (k, r) could be different.

One option to merge the transcripts of the two models is using a word-level aligner like Rover [13]. However, for accent errors we expect the service to be wrong only on a sub-part of a word, say a ‘t’ being wrongly identified as a ‘d’. The local transcript \mathbf{c} might correct some accent errors while missing out on other parts of the word. In general, the local model is expected to be weaker than the service on all but the accent errors, for the client to want to pay for the service. As an example, consider a gold transcript `toasted` that gets recognized as `posted` and `to state` by the service and local models, respectively. The service model fails to recognize the `t` in `toasted` and outputs `posted`. The local model recognizes the `t` but yields `to state`. To reconstruct the correct word in such cases we need a finer-grained splicing at the sub-word level.

Given the prevalence of character-level models in modern ASR systems we then sought to splice the two transcripts at

1	S_t	-	p	o	-	s	t	e	d	d
	$P_t(S_t)$	6e-5	1e-11	1	0.01	0.93	0.99	0.44	0.29	0.98
2	d_t	t	t	o	-	s	t	a	t	d
	$P_t(d_t)$	0.99	0.99	1.0	0.98	0.93	0.99	0.55	0.64	0.98
3	r_t	t	t	o	-	s	t	e	d	d
	$P_t^s(r_t)$	0.62	0.99	1.0	0.61	0.93	0.99	0.66	0.57	0.98
	$P_t(r_t)$	0.99	0.99	1.0	0.98	0.93	0.99	0.44	0.29	0.98
	Frame t	1	2	3	4	5	6	7	8	9

Table 1: *Example: Client model revising frame-level character distribution $\mathbf{P} \rightarrow \mathbf{P}^s$ using service transcript $\mathbf{s} = \text{'posted'}$ in FineMerge. $d_t = \text{argmax}_c P_t(c)$ and $r_t = \text{argmax}_c P_t^s(c)$. First row shows aligned service characters and their probability from \mathbf{P} , second row shows the modes of the \mathbf{P} distribution, third row shows the $\text{argmax}_c r_t$ of the revised distribution and its probability from the revised and original distribution.*

the character-level. Designing a good character-level merging strategy is challenging because of large divergences between the two outputs, both because of the differential strengths of their acoustic models and the introduction of unheard characters when biasing with their respective language models. Strategies like combining the characters from the two outputs using Rover-like algorithms fail to distinguish between the two types of errors in the absence of accurate character-level confidence from the service. For example, aligning the characters in `posted` with `to state` yielded `too sttd`

We finally designed an algorithm that exploits white-box access to the local model \mathcal{C} to guide its decoding using the service transcript \mathbf{s} , instead of merging a fixed \mathbf{c} from \mathcal{C} .

We assume the local model \mathcal{C} is trained using the standard CTC loss invoked on frame-level character distributions [22] that maximizes likelihood of the target \mathbf{y} by marginalizing over all alignments compatible with \mathbf{y} . During inference, the trained model generates the distribution over alignments for an input \mathbf{x} and predicts character distributions $P_1, \dots, P_T \triangleq \mathbf{P}$ at each of the T frames of the input. From these probability distributions, an output sequence \mathbf{c} is recovered using beam-decoding in conjunction with a language model (LM).

We guide this inference using the service transcript in two steps: First align the service characters with each frame of the local model using its frame-level probability distributions \mathbf{P} . Next revise \mathbf{P} to selectively support \mathbf{s} . We elaborate these steps next. A pseudo code appears in Algorithm 1.

Aligning service characters Our first step is to expand out the characters in \mathbf{s} over the T frames by repeating characters or inserting blanks so as to maximize the probability of the aligned characters as per P_1, \dots, P_T . Let S denote the highest probability expanded character sequence. An example is shown in Table 1 where $\mathbf{s} = \text{posted}$ is aligned over $T = 9$ frames and the resulting S is shown in the first row. The full \mathbf{P} cannot be shown but we show the probability of the aligned character below it and the maximizing character probability in the second row. Such a forced alignment of \mathbf{s} with \mathbf{P} can be solved optimally using a simple Viterbi-like dynamic programming algorithm. The algorithm processes \mathbf{s} time-synchronously over the T frames such that either a symbol from \mathbf{s} or a blank is produced as output at each frame. This is referred to as ‘‘Viterbi-align’’ in Algorithm 1. Successfully aligning the service characters requires an additional consideration. The server’s output \mathbf{s} contains characters that can be attributed to both accent errors and cascaded language model errors. We therefore smooth \mathbf{P} distribution by adding a small constant 10^{-20} to all probability

Algo 1: The FineMerge Inference Algorithm

Input: \mathbf{x} : Input audio with T frames
 \mathcal{C} : Local model fine-tuned on target accent
 ψ : service probability threshold
 ω : service weight for mixing
 γ : probability of blank

Output: Final transcript

```
1  $\mathbf{s}, \mathbf{p} \leftarrow$  Transcript, token-confidence from Service on  $\mathbf{x}$ 
2  $P_1, \dots, P_T \leftarrow$  Frame-level probability from  $\mathcal{C}(\mathbf{x})$ 
3  $S_1, \dots, S_T \leftarrow$  Viterbi-align( $\mathbf{s}$ , Smooth( $\mathbf{P}$ ))
4 for  $t \leftarrow 1$  to  $T$  do
5   if  $\psi < P_t[S_t] < \max_c P_t[c]$  then
6      $\omega_t \leftarrow \gamma$  if  $S_t$  is blank else  $\omega \mathbf{p}$ 
7     [word index of  $S_t$ ]
8      $P_t^s \leftarrow (1 - \omega_t)P_t + \omega_t \text{oneHot}(S_t)$ 
9   else
10     $P_t^s \leftarrow P_t$ 
11 return Beam-decode using  $\mathbf{P}^s$  and local LM of  $\mathcal{C}$ 
```

entries so even unheard characters get non-zero probability.

Revising \mathbf{P} with \mathbf{s}, \mathbf{p} Now each frame t is aligned with a character S_t in service. We need to revise \mathbf{P} so as to 'support' the aligned characters of service while ignoring those characters which may have been erroneously introduced during LM-based decoding. For this, we boost the probability of the aligned service characters in P_t on those frames t where the probability $P_t(S_t)$ is less than the maximum probability in P_t but greater than a threshold ψ . The lower limit ψ is to suppress those characters in \mathbf{s} which are not 'heard' at all by the client's acoustic model, and are likely to have been introduced by the LM. The amount of boosting is product of a hyper-parameter ω and the confidence of the parent word of S_t . If S_t is blank, we use a fixed probability γ . We use \mathbf{P}^s to denote the \mathbf{P} distribution after this revision with \mathbf{s} . In Table 1, we show the mode of the revised distribution \mathbf{P}^s in row 3. Note, how 'p' in frame 2 was ignored in favor of the gold character 't' since $P_2(p)$ has a very small probability (1e-11). In frame 7, the 'e' from the service was used to boost the probability of 'e' in the P_8 distribution from 0.44 to 0.66. Likewise in frame 8. Greedy decoding on the revised distribution yields `to sted` which is closer to the gold token `toasted` than either the service token `posted` or the local token `to state`. Beam-decoding on the revised \mathbf{P}^s recovers the gold token.

The above merging algorithm is simple and requires tuning only three hyper-parameters. Since client's labeled data is limited, we found that more elaborate attention-based merging models using several parameters did not perform well.

4. Experiments

We evaluate FineMerge on two accents and three service combinations and contrast against four other methods. We present anecdotes and analyze the accent adaptations we achieve.¹

Datasets We used the Mozilla Common Voice v4 (MCV-v4) dataset. The dataset is crowd-sourced and contains 1,118 hours of validated speech data of varying accents. We extracted

¹Our code is available at <https://github.com/Kartik14/FineMerge>.

Method	WER		CER	
	Ind	Aus	Ind	Aus
Local	27.99	24.41	16.98	14.55
Service	22.32	23.52	11.96	13.27
Rover	21.12	18.04	11.95	9.81
LM rescore	22.10	23.42	12.10	13.56
FineMerge	18.45	16.90	10.65	9.33

Table 2: Overall comparison on WER and CER for Indian and Australian accented data.

around 28K Indian and 27K Australian accented utterances, amounting to 37 and 35 hours of speech, respectively. For each accent, we split into train, validation and test sets roughly in the ratio 85-5-10 ensuring no overlap among speakers and transcripts. The MCV-v4 audio clips were normalized in a pre-processing step.

Service and Local Models We used Google Cloud Speech to Text API [1] as our default service model, and include two other service models later. For the local, we used the DeepSpeech2 (DS2) [11] model pretrained on the LibriSpeech corpus [12] and then fine-tuned individually for each accent. We used a trigram LM trained on sentences from the MCV-v4 corpus after removing sentences overlapping with test sets. DS2 parameters α (for LM weight) and β (to encourage more words) were also fine-tuned on the validation set for each accent. The hyper-parameters of our method ω, ψ, γ were also tuned on the validation set for each accent.

Methods compared We measure word error rates (WER) on five different models: the service model, the local model, Rover [13] on the confidence weighted transcripts of service and local model, LM rescoring top-N whole transcripts from service, and our FineMerge method.

Overall Results In Table 2 we show the WERs on the Indian and Australian accents for these five methods. Observe that overall the error rate of Service is lower than that of accent-adapted Local. Rover's word-level merging provides significantly improved results than either of the two indicating that the two models exhibit complementary strengths. LM rescoring does not improve results much, establishing that the local LM may not have much impact on the improved results. Our algorithm FineMerge provides the greatest gains in WER over all methods. For the Australian accent, we obtain a 28% relative reduction in WER over either of the service and client models. Table 3 presents some anecdotes which show how the fine-grained merging enables us to recover the highlighted word, even when neither the service nor client models contain that word.

Comparing methods of character alignment A centerpiece of our method is Viterbi aligning \mathbf{s} with the frame-level character probability distribution. We show that this achieves a character-level alignment that is more accurate than existing methods by focusing only on character error rate (CER) before beam-decoding. The last two columns in Table 2 presents CER of Local (before LM decoding), Service (as is), Rover applied at the character-level on these two, LM rescoring, and FineMerge's after selecting the modes of the revised distribution \mathbf{P}^s i.e., before LM decoding. We observe that FineMerge's CER is much lower particularly for Indian accent. This explains that

	INDIAN	AUSTRALIAN
Gold	everyone toasted the ..	nora finds herself ugly ..
Service	everyone posted the ..	nora van to self ugly ..
Local	everyone to state the ..	nor iphones herself ugly ..
Rover	everyone to posted the ..	nor to self ugly ..
FineMerge	everyone toasted the ..	nora finds herself ugly ..
Gold	for a brief time ..	hannelore is an ..
Service	soda beef time ..	i don't know what is an ..
Local	for a breese time ..	hailar is an ..
Rover	for a beef time ..	i don't know what is an ..
FineMerge	for a brief time ..	hannelore is an ..
Gold	the condition also occurs..	..rope a bull while on a
Service	definition of circus..	..work a bowl while on a
Local	the condition also acres..	..rope the ball while on a
Rover	the definition also circus..	..work a bowl while on a
FineMerge	the condition also occurs	..rope a bull while on a

Table 3: Anecdotes comparing transcripts of Indian and Australian accents speech from five different methods.

the main reasons for our gains is due to our novel frame-level fine-grained merging algorithm.

Varying Quality of Service Model In addition to the default Google Speech API service (G-US), we evaluate on two other models as service — a second Google speech-to-text model (G-Video) [23] meant for transcribing audio of video files, which works significantly better for MCV-v4 utterances because of their low-fidelity, and Jasper [24], a recent end-to-end convolutional neural ASR model trained on the LibriSpeech dataset. We note here that we opted for G-US rather than Google’s ASR API for Indian English because of the latter’s poor performance (compared to G-US) on MCV-v4 utterances that are low bandwidth. Table 4 shows the results. WER of local stays the same since service has no role during its training. We see a wide difference in accuracies across the service models. G-Video is the most accurate, but even in this case FineMerge is able to obtain a relative WER reduction by at least 3%. The Jasper model is worse than local Indian fine-tuned, yet FineMerge achieves more than 15% relative WER reduction over both Service and Local. This shows that the hyperparameters of our service-guided local inference adapts even to a weaker service model.

Method	Indian			Austrian		
	G-US	G-Video	Jasper	G-US	G-Video	Jasper
Local		27.99			24.41	
Service	22.32	13.77	31.82	23.52	11.08	19.56
Rover	21.12	20.51	26.95	18.04	13.84	17.57
LM rescore	22.10	13.37	31.38	23.42	10.99	19.35
FineMerge	18.45	13.36	23.72	16.90	10.68	16.07

Table 4: Effect of changing service model

Importance of Accent Adaptation One question is if our gains are merely due to ensembling of any two independent models adapted to test data domain, or did we specifically adapt accent. To answer this, we run FineMerge with a local model fine-tuned on a similarly-sized MCV corpus from a different accent. Table 5 compares our WER to the WER on client models fine-tuned on US accented samples. Observe that FineMerge outperforms service even when the local model is fine-tuned on a different accent. This captures the base benefit of ensembling. However, after fine-tuning on data of its own accent the

Test accent	Service	FineMerge with	
		(Indian/Aus)-Local	US-local
Indian	22.32	18.45	21.01
Aus	23.52	16.90	20.66

Table 5: WER comparison with different local models.

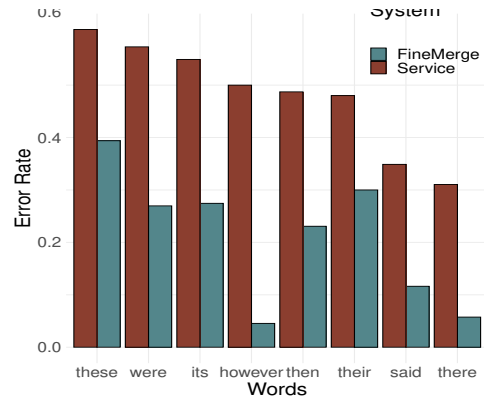


Figure 1: Highest reductions in error per word on Indian-accented test samples.

gains are higher. For Aus accent, service WER of 23.52 drops to 20.66 with FineMerge on Indian-Local but drops further to 16.90 on Aus-Local.

Figure 1 shows the largest reductions in errors per word on Indian test samples obtained by FineMerge over service. Error rates are cut in half for most words revealing FineMerge’s ability to do accent adaptation. Word “however” is an interesting example to highlight. The diphthong /AW/ in “however” has a range of phonetic realizations across Indian speakers that has been investigated in prior work [25]. This variability is difficult for the service to accurately model, while FineMerge cuts the errors on “however” down to 5% from 50%. Another interesting example is “were”. The phonemes /v/ and /w/ are indistinguishable in most Indian languages, making minimal pairs like *veil* and *wail* homophones when articulated by Indian speakers.

/DH/-initial words like “then”, “these”, “their” and “there” are other likely targets of accent errors due to the lack of dental fricatives like /DH/ in most Indian languages. FineMerge is able to substantially reduce these errors.

5. Conclusion and Future Work

In this paper we motivated and introduced the problem of black-box adaptation of an ASR service. We presented a novel coupling of an open-source accent adapted model with the black-box service model to fix accent errors in an otherwise strong service model. We presented FineMerge an algorithm that achieves a fine-grained mixing of the service output and local frame-level distributions. We show that such fine-grained mixing is specifically effective in fixing accent errors that word-level mixing cannot fix. Our strategy achieves upto 28% reduction in word-error rate over service APIs of varying grades of quality. Future work could consider combining outputs from multiple services and fixing both dialect and accent differences.

6. Acknowledgements

This research was partly sponsored by IBM AI Horizon Networks - IIT Bombay initiative.

7. References

- [1] "Cloud speech-to-text api." [Online]. Available: <https://cloud.google.com/speech-to-text/docs/reference/rest>
- [2] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Toups, J. R. Rickford, D. Jurafsky, and S. Goel, "Racial disparities in automated speech recognition," *Proceedings of the National Academy of Sciences*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [3] Y. Huang, D. Yu, C. Liu, and Y. Gong, "Multi-accent deep neural network acoustic model with accent-specific top layer using the kld-regularized model adaptation," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [4] M. Chen, Z. Yang, J. Liang, Y. Li, and W. Liu, "Improving deep neural networks based multi-accent mandarin speech recognition using i-vectors and accent-specific top layer," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [5] X. Yang, K. Audhkhasi, A. Rosenberg, S. Thomas, B. Ramabhadran, and M. Hasegawa-Johnson, "Joint modeling of accents and acoustics for multi-accent speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1–5.
- [6] A. Jain, M. Upreti, and P. Jyothi, "Improved accented speech recognition using accent embeddings and multi-task learning," in *Proceedings of Interspeech*, 2018.
- [7] T. Viglino, P. Motlicek, and M. Cernak, "End-to-end accented speech recognition," 2019.
- [8] K. Rao and H. Sak, "Multi-accent speech recognition with hierarchical grapheme based models," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4815–4819.
- [9] E. K. Ringger and J. F. Allen, "Error correction via a post-processor for continuous speech recognition," in *IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, vol. 1. IEEE, 1996, pp. 427–430.
- [10] R. Corona, J. Thomason, and R. Mooney, "Improving black-box speech recognition using semantic parsing," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 2017, pp. 122–127.
- [11] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International conference on machine learning*, 2016, pp. 173–182.
- [12] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [13] J. G. Fiscus, "A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover)," in *IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*. IEEE, 1997, pp. 347–354.
- [14] V. Soto, O. Siohan, M. Elfeky, and P. Moreno, "Selection and combination of hypotheses for dialectal speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5845–5849.
- [15] D. Vergyri, L. Lamel, and J.-L. Gauvain, "Automatic speech recognition of multiple accented english data," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [16] S. Sun, C.-F. Yeh, M.-Y. Hwang, M. Ostendorf, and L. Xie, "Domain adversarial training for accented speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4854–4858.
- [17] Y.-C. Chen, Z. Yang, C.-F. Yeh, M. Jain, and M. L. Seltzer, "Aipnet: Generative adversarial pre-training of accent-invariant networks for end-to-end speech recognition," *arXiv preprint arXiv:1911.11935*, 2019.
- [18] J. J. Humphries, P. C. Woodland, and D. Pearce, "Using accent-specific pronunciation modelling for robust speech recognition," in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP'96*, vol. 4. IEEE, 1996, pp. 2324–2327.
- [19] Y. Zheng, R. Sproat, L. Gu, I. Shafran, H. Zhou, Y. Su, D. Jurafsky, R. Starr, and S.-Y. Yoon, "Accent detection and speech recognition for shanghai-accented mandarin," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [20] S. Watanabe and J. Le Roux, "Black box optimization for automatic speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 3256–3260.
- [21] A. Kastanos, A. Ragni, and M. Gales, "Confidence estimation for black box automatic speech recognition systems using lattice recurrent neural networks," *arXiv preprint arXiv:1910.11933*, 2019.
- [22] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [23] "Cloud speech-to-text api video model." [Online]. Available: <https://cloud.google.com/speech-to-text/docs/transcription-model>
- [24] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. T. Gadde, "Jasper: An end-to-end convolutional neural acoustic model," *arXiv preprint arXiv:1904.03288*, 2019.
- [25] O. Maxwell and J. Fletcher, "The acoustic characteristics of diphthongs in indian english," *World Englishes*, vol. 29, no. 1, pp. 27–44, 2010.