

CANTONESE AUTOMATIC SPEECH RECOGNITION USING TRANSFER LEARNING FROM MANDARIN

Bryan Li¹, Xinyue Wang¹

Columbia University¹
New York, NY
[bl2557, xw2368]@columbia.edu

Homayoon Beigi^{1,2}

Recognition Technologies, Inc.²
White Plains, NY
beigi@recotechnologies.com

ABSTRACT

We propose a system to develop a basic automatic speech recognizer (ASR) for Cantonese, a low-resource language, through transfer learning of Mandarin, a high-resource language. We take a time-delayed neural network trained on Mandarin, and perform weight transfer of several layers to a newly initialized model for Cantonese. We experiment with the number of layers transferred, their learning rates, and pre-training i-vectors. Key findings are that this approach allows for quicker training time with less data. We find that for every epoch, log-probability is smaller for transfer learning models compared to a Cantonese-only model. The transfer learning models show slight improvement in CER.

Index Terms— automatic speech recognition, transfer learning, time-delayed neural networks, Cantonese ASR, low-resource languages

1. INTRODUCTION

Much of the work done thus far on automatic speech recognition has focused on a subset of high-resource languages, such as English, Spanish and Mandarin Chinese – these languages are fortunate to have expert-built dictionaries, part-of-speech taggers, language models, as well as large-amount of high-quality voices recorded in professional studios. All these allow for development of state-of-art ASR systems. However, there are approximately 6500 languages in the world, many of which spoken by millions of people, yet have no such resources and are little studied for speech recognition. Interest has recently been taken in working with low-resource languages. This work both generalizes and strengthens language and speech modeling technologies, and opens up speech technologies to a new cohort of speakers.

Cantonese is a variety of Yue Chinese spoken by over 100 million people in Southeastern China. Most speakers of Cantonese are fluent in Mandarin, the national language, and one at the forefront of speech research. As a result, there is a general lack of interest in Cantonese such that it is a low-resource language. This is further complicated by its lack of

a standardized written form¹ because Cantonese is a spoken-dialect, which has made it difficult for quality textual and audio corpora to be collected.

Though varieties of Chinese are often thought of as merely dialects, they are not mutually intelligible and just as distinct as family like Romance languages. Some key differences between Cantonese and Mandarin are: 6 vs. 4 tones, presence of final stops, lack of initial retroflexes, and colloquial lexicon. Still, the languages share a common ancestor and many similarities. Thus, we aim to leverage the myriad of resources available in Mandarin Chinese to train a state-of-the-art Mandarin ASR system. From here, we apply transfer learning to initialize parameters of a Cantonese ASR system, training further on a limited Cantonese dataset.

2. RELATED WORK

2.1. Transfer Learning

Transfer learning is a vital technique that generalizes models trained for one task to other tasks [1]. In speech processing, some common patterns represented by features such as MFCCs and pitches are shared across languages, especially for similar/geographically closely located languages such as Cantonese and Mandarin. Neural network models have many parameters, and as such depend on large volumes of data to learn the patterns of a problem. Furthermore, they are not robust to low-quality training data with errors. Both issues are prevalent when working with low-resource languages. Transfer learning helps to alleviate these issues, since we are able to use well-maintained corpora from high-resource languages. A key idea is that the features learned by deep neural network models are more language-independent in earlier layers, and more language-dependent in later layers [1].

Zoph et al. [2] show the effectiveness of applying transfer learning on different languages in the domain of machine translation. Initializing from parameters of a parent model in French-English and training child models in low-resource languages such as Turkish and Urdu, they see significant in-

¹Hong Kong Cantonese is, but we consider mainland Cantonese only.

creases in BLEU scores. They find better gains when applying transfer learning to related languages, such as French and Spanish, to unrelated languages, such as French and German.

Kunze et al. [3] develop a German ASR system using the transfer learning approach of *model adaptation* [1]. This involves first training a model on one language (in this case English), and then retrain all or parts of it on a different language (German). They use an 11-layer convolutional neural network (CNN), and experiment with fixing k layers. They find that such weight initialization allows for much faster training without a decrease in quality. Freezing more layers (higher k) decreases training time, but the best-performing model allows all layers to learn ($k=0$). Our report seeks to experiment similarly, but with a time-delayed neural network (TDNN) on Mandarin to Cantonese.

2.2. State-of-Art for Cantonese ASR

We primarily consider the BABEL Cantonese dataset, further described in Section 3. Karafiat et al. [4] explore feature extraction using a 6-layer stacked bottleneck neural network. These outputs are used as features for a GMM-HMM recognition system. They concatenate these features with fundamental frequency f_0 , reasoning that pitch is important to Cantonese, a tonal language. Further, they apply region-dependent transforms, a speech-specific GMM optimization method [4], to all these features, and arrive at a character-error rate (CER) of 42.4%.

Further work is done by Du et al. [5], who note an inherent difficulty: because there is no standardized form of Cantonese, it is difficult to perform web-scraping tasks in isolating Cantonese text. Therefore, they combine the BABEL Cantonese transcriptions with data-augmented Cantonese text machine translated from transcriptions of Mandarin telephone conversations. Using 88 features including pitch, Mel-PLP, and TRAP-DCT input into a bottleneck DNN, they are able to improve the CER to 40.2%.

3. DATASETS USED

The Cantonese dataset is IARPA Babel Cantonese [6]. This includes 215 hours of Cantonese conversational and scripted telephone speech, along with corresponding transcripts. Speakers represent dialects from the Chinese provinces of Guangdong and Guangxi. The 952 speakers range from ages 16 to 67, and the male-female ratio is 48%:52%. Audio was collected from mobile and landline telephones, and from environments such as the street, at home, and inside a vehicle. We train and test on only the conversational data (75% of all data), as done by all prior work with this dataset.

The Mandarin dataset is AISHELL2 [7], which consists of 1000 hours of clean read-speech data [3], along with corresponding transcripts. The 1991 speakers range from ages 11 to 40 that speak on 12 topics such as voice control, news, and

technology. The male-female ratio is 40%:60%. Audio was collected from three sources. From left to right: iOS device, a high-fidelity microphone, and an Android device. We use only the iOS-recorded data, as it produced the best results for the Mandarin model.

3.1. Preprocessing

The Mandarin data has a sample rate of 16 kHz, whereas the Cantonese data has a sample rate of 8 kHz. To be able to work with the same extracted features, we downsample Mandarin audio to 16 kHz. However, the audio quality of BABEL is still worse, and about half of the audio is silence.

All data preprocessing, including resampling and conversion to wav format, was done with `sox`. Speed perturbation of each file (90%, 100%, 110%) is performed to augment the training data, and volume perturbation is performed to make models more invariant to test data volume.

3.2. Feature Extraction

The GMM models are trained with 13-dimensional Mel-frequency cepstral coefficients (MFCCs) and 3 pitch features. Note that Babel originally used Perceptual Linear Prediction (PLP) features. The TDNN models are trained with 39-dimensional MFCCs and 4 pitch features. They also include for each frame-wise input a 100-dimensional i-vector. Two extractors are trained, one for each of Cantonese and Mandarin. We extract Cantonese i-vectors using one or the other for different models.

4. MODEL ARCHITECTURE

Our model is implemented in Kaldi, and follows a fairly standard Kaldi pipeline. Our recipe is written with reference to existing ones –TEDLIUM, BABEL, AISHELL2.

4.1. Language Model

Written Chinese does not have spaces between words, so ASR systems require sophisticated word segmentation. AISHELL2 [7] includes an open-source dictionary called DaCiDian, which maps Chinese words in two layers, first from word to PinYin syllables [8], and second from PinYin to phoneme. Word segmentation is done with Jieba [9], which uses a prefix-tree based search approach, and the Viterbi algorithm for out-of-vocabulary (OOV) words.

In parallel with AISHELL, we present a Cantonese dictionary called DaaCiDin, which maps Cantonese words to Yale romanization system syllables, then maps these to X-SAMPA phonemes. Its vocabulary is compiled from training set transcripts of Babel Cantonese.

The statistical language model is implemented in SRILM, and consists of n-grams of size 2, 3, and 4 using both Kneser-

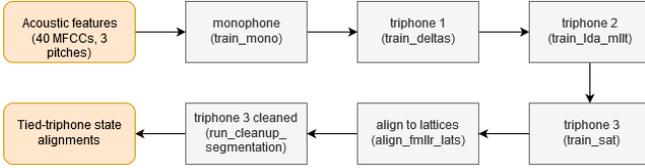


Fig. 1. GMM-HMM acoustic model

Ney and Good-Turing smoothing. Silence and OOV words are also accounted for.

4.2. Acoustic Model

The acoustic model consists of two stages: a GMM-HMM model, and a TDNN-LSTM model.

4.2.1. GMM

The Gaussian Mixture Model is shown in Figure 1. It is trained with an input dimension of 16 (13 MFCCs plus 3 pitches). A monophone model is trained as the starting point for the triphone models. Then, a larger triphone model is trained using delta features. To replace the delta features, linear discriminant analysis (LDA) is applied on stack of frames, and MLLT-based global transform is estimated iteratively. Next comes the Speaker Adaptive Training (SAT) stage, whose outputs are aligned to lattices. Finally, we run Kaldi clean-up scripts so all audio in the training data matches the transcripts. The GMM model outputs the tied-triphone state alignments.

4.2.2. TDNN

The TDNN-LSTM model consists of time-delayed neural network [10] and long-short term memory layers. It alternates between the layer types as shown in Figure 2.

We first train a TDNN on AISHELL2. The objective function is lattice-free maximum mutual information (LF-MMI) [11]. The network also uses sub-sampling as described in [10] and is not fully connected – the input of each hidden layer is a frame-wise spliced output of its preceding layer [7]. The input to the network is high resolution MFCC with cepstral normalization plus pitches, making the dimension 43. In addition, a 100-dimensional i-vector is attached to each frame-wise input using the extractor trained on AISHELL2 to encode speaker dependent information. i-vectors use as features high-resolution MFCCs, pitches, and the GMM-HMM tied-triphone state alignments.

In the transfer learning stage, we adapt the TDNN trained for Mandarin as described above to the Cantonese language. We initialize the weights of the first k layers of the TDNN to be those of the Mandarin model, and set their learning rate = x . Then we randomly initialize the weights of remaining layers,

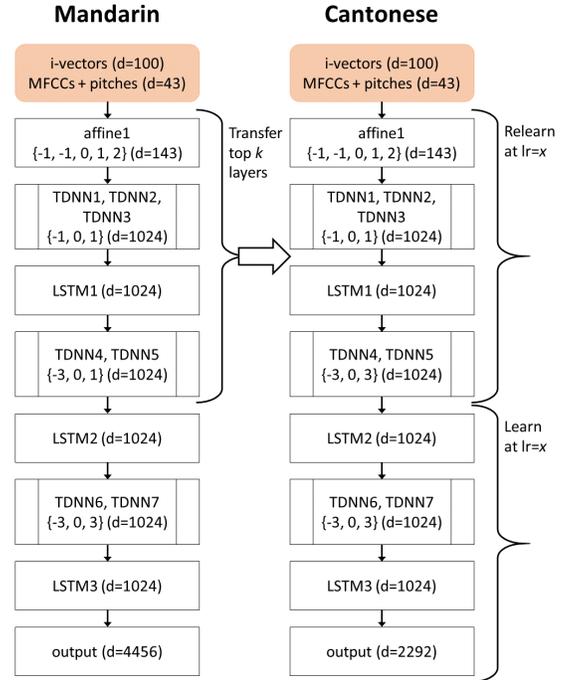


Fig. 2. TDNN-LSTM model + transfer learning from Mandarin to Cantonese

and set their learning rate = 1.0. i-vectors are obtained using either the AISHELL2 or BABEL Cantonese i-vector extractor depending on our experimental configuration.

5. EXPERIMENTS

We perform several experiments to explore the space of transfer learning. All our experiments use the model architecture as described above, and modify the learning rate of the transferred layers when training the Cantonese model. The learning rates we used are 0, 0.25, and 1.0. These are relative to the learning rate of the final block, which is always 1.0. $lr=0$ corresponds to fixing transferred learning weights, whereas 0.25 and 1.0 allow the model to adjust them as necessary.

For $lr=0.25$, we additionally train a model that uses the i-vector extractor of AISHELL2 on the BABEL Cantonese dev set. For comparison purposes, we also train two baseline Cantonese models: one using the TDNN recipe from BABEL, and one using the AISHELL2 architecture without transfer learning (all weights initialized to zero).

The same hyperparameters are used to train all models: number of epochs=4, initial learning rate=0.001, final learning rate=0.0001, minibatch size=128, cross-entropy regularization=0.1, and frames per example=150, 110, and 90. The dropout schedule is 0, 30% after 50% of the data is seen, and 0 for the last layer.

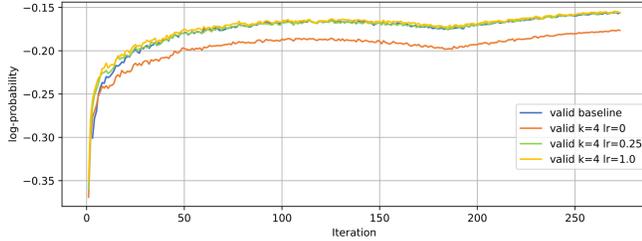


Fig. 3. Plot of log-probability vs iteration for output

6. RESULTS AND DISCUSSION

In this section we report experimental results of transfer learning. First, we compare model performance in terms of log-probability and training time. Second, we compare ASR performance in terms of character and word error rates.

6.1. Metrics Used

Word error rate (WER) is the standard metric used for speech recognition systems, and is given by:

$$WER = \frac{S + D + I}{N}$$

S, D, and I are the number of substitutions, deletions, and insertions respectively. N is the total number of words in the reference transcription. WER ranges from 0-300%.

Character error rate (CER) follows the same formula, but with characters as the unit. In written Chinese, the majority of words are made of two characters. Following prior work, we primarily report CER results.

6.2. Model Loss

Figure 3 compares the the log-probability (log-prob) over time for different models. The results show that log-prob (of cross-entropy loss) is always higher for transfer learned models, for both train and validation sets. Higher learning rates correspond to higher log-prob.

These findings are consistent with those of Kunze et al. [3]. Because the languages are not the same (for example, Cantonese uses long vowels instead of diphthongs [12, 13]), it is better to allow the earlier layer weights to adjust themselves than to fix them.

6.3. Reduced Computing Time

Given that languages share common features, earlier layers in the model should contain common features that can be transferred to that of a different language. Thus, we can fix the earlier layers of the original Mandarin model, and this decreases the average training time per iteration because the network will need to back-propagate through fewer layers. Train times

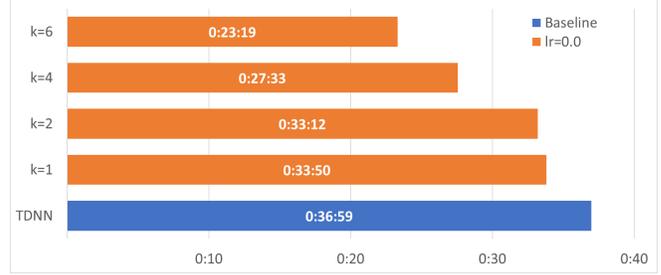


Fig. 4. Comparison of training time in minutes per iterations for different learning rates

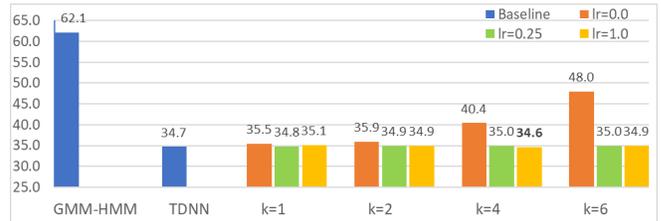


Fig. 5. Character Error Rate of transfer-learned models vs baseline

depending on k are shown in Figure 4. Note that for all models $lr=0$, train time is equivalent to the baseline train time.

However, from our discussion in section 6.2, loss is always smaller at any point for non-zero learning rate. Thus, we figure that in terms of achieving small loss, there is no reason to freeze the early layers. Nevertheless, when we compare the log-probability of the transferred model with the baseline Cantonese model trained from scratch, we see the transferred model starts off with a much lower loss and thus is able to achieve the same loss with shorter time. Therefore, it is beneficial in terms of computing time to transfer learn with a good weight initialization.

6.4. Character and Word Error Rates

Figure 5 shows the speech recognition performance of the various systems on the development set. For $lr=0.0$, the more transferred layers, the worse results are – this makes sense because of the distance between the languages. We should still do some learning. For $lr=0.25$ and 1.0 , results are similar across the board. We see that our best model, $k=4$ and $lr=1.0$, gives a slight increase of 0.1% CER vs the baseline.

We do not report results for the experiments with pre-training i-vectors on Mandarin, and transferring to Cantonese, as they do not improve performance. We suspect that weight transfer could help here, as with the TDNN-LSTM, but leave that to future work.

7. FUTURE WORK

For future work, we have several approaches in mind. First, despite downsampling AISHELL-2 to 8 kHz to match BABEL Cantonese, the latter still has far worse audio quality. Therefore, we will try adding acoustic noise Mandarin, and training the model to be transferred on that. Second, we aim to explore the language model more. The LMs are only trained on acoustic transcripts, and we can add more data. Also, prior linguistic work has shown that correspondence can be made between Cantonese and Mandarin phonemes (and to a lesser extent, tones), so we can create a mapping between phonemes. Most written Chinese characters are words in both Cantonese and Mandarin, so that would help with mapping. Finally, fine-tuning the AISHELL-2 i-vectors after pre-training would be a good experiment.

8. CONCLUSION

We propose a system to develop a basic ASR for Cantonese through transfer learning of Mandarin. We first trained a TDNN for Mandarin using large amounts of data from AISHELL2, and adapted the model to Cantonese by transferring the weights and adjusting the relative-learning rates of earlier and later layers. We conducted multiple experiments setting $lr = 0, 0.25$ and 1 and show that log-probability is always higher for transfer learned models. In addition, we showed reduced computation time for freezing earlier.

We achieve a new state-of-the-art on BABEL Cantonese (CER=34.6%) by transferring up to TDNN4 and using $lr=1.0$. Results are not as promising as prior work has been, but we believe our work is still as step in the right direction. We hope to make the transfer learning methods more informed than simple model adaptation.

9. REFERENCES

- [1] D. Wang and T. F. Zheng, "Transfer learning for speech and language processing," *CoRR*, vol. abs/1511.06066, 2015. [Online]. Available: <http://arxiv.org/abs/1511.06066>
- [2] T. Kocmi and O. Bojar, "Trivial transfer learning for low-resource neural machine translation," *CoRR*, vol. abs/1809.00357, 2018. [Online]. Available: <http://arxiv.org/abs/1809.00357>
- [3] J. Kunze, L. Kirsch, I. Kurenkov, A. Krug, J. Johannsmeier, and S. Stober, "Transfer learning for speech recognition on a budget," in *Proceedings of the 2nd Workshop on Representation Learning for NLP*, 2017, pp. 168–177.
- [4] K. Martin, F. Grzl, M. Hannemann, V. Karel, and . Jan, "But babel system for spontaneous cantonese," 09 2013.
- [5] G. Huang, A. Gorin, J. Gauvain, and L. Lamel, "Machine translation based data augmentation for cantonese keyword spotting," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 6020–6024.
- [6] T. Andrus, E. Dubinski, J. G. Fiscus, B. Gillies, M. Harper, T. J. Hazen, B. Hefright, A. Jarrett, W. Lin, J. Ray, A. Rytting, W. Shen, E. Tzoukermann, and J. Wong, "Iarpa babel cantonese language pack iarpa-babel101b-v0.4c ldc2016s02," 2016, web Download. [Online]. Available: <https://catalog.ldc.upenn.edu/LDC2016S02>
- [7] J. Du, X. Na, X. Liu, and H. Bu, "AISHELL-2: transforming mandarin ASR research into industrial scale," *CoRR*, vol. abs/1808.10583, 2018. [Online]. Available: <http://arxiv.org/abs/1808.10583>
- [8] S. Duanmu, *The phonology of standard Chinese*. Oxford University Press, 2007.
- [9] J. Sun. (2012) Jieba chinese word segmentation tool. [Online]. Available: <https://github.com/fxsjy/jieba>
- [10] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *INTERSPEECH*, 2015.
- [11] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi." in *Interspeech*, 2016, pp. 2751–2755.
- [12] F. Seide and N. J. Wang. (1998) Phonetic modelling in the philips chinese continuous-speech recognition system. [Online]. Available: <http://w.colips.org/conferences/isclp2006/anthology/1998/Papers/ASR-A3.pdf>
- [13] J. Mar. (2002) Sampa for cantonese. [Online]. Available: <https://www.phon.ucl.ac.uk/home/sampa/cantonese.htm>