



# Speech Model Pre-training for End-to-End Spoken Language Understanding

Loren Lugosch<sup>1</sup>, Mirco Ravanelli<sup>1</sup>, Patrick Ignoto<sup>2</sup>, Vikrant Singh Tomar<sup>2</sup>, Yoshua Bengio<sup>1</sup>

<sup>1</sup>Université de Montréal / Mila, Canada

<sup>2</sup>Fluent.ai, Canada

{lugoschl, mirco.ravanelli, yoshua.bengio}@mila.quebec  
{patrick.ignoto, vikrant}@fluent.ai

## Abstract

Whereas conventional spoken language understanding (SLU) systems map speech to text, and then text to intent, end-to-end SLU systems map speech directly to intent through a single trainable model. Achieving high accuracy with these end-to-end models without a large amount of training data is difficult. We propose a method to reduce the data requirements of end-to-end SLU in which the model is first pre-trained to predict words and phonemes, thus learning good features for SLU. We introduce a new SLU dataset, Fluent Speech Commands, and show that our method improves performance both when the full dataset is used for training and when only a small subset is used. We also describe preliminary experiments to gauge the model's ability to generalize to new phrases not heard during training.

**Index Terms:** speech recognition, spoken language understanding, end-to-end models, transfer learning

## 1. Introduction

Spoken language understanding (SLU) systems infer the *meaning* or *intent* of a spoken utterance [1]. This is crucial for voice user interfaces, in which the speaker's utterance needs to be converted into an action or query. For example, for a voice-controlled coffee machine, an utterance like "make me a large coffee with two milks and a sugar, please" might have an intent representation like `{drink: "coffee", size: "large", additions: [{type: "milk", count: 2}, {type: "sugar", count: 1}]}`.

The conventional SLU pipeline is composed of two modules: an automatic speech recognition (ASR) module that maps the speech to a text transcript, and a natural language understanding (NLU) module that maps the text transcript to the speaker's intent [2–4]. An alternative approach that is beginning to gain popularity is end-to-end SLU [5–9]. In end-to-end SLU, a single trainable model maps the speech audio directly to the speaker's intent without explicitly producing a text transcript (Fig. 1). Unlike the conventional SLU pipeline, end-to-end SLU:

- directly optimizes the metric of interest (intent recognition accuracy),
- does not waste modeling effort on estimating the text, thus yielding a more compact model and avoiding an error-prone intermediate step involving search algorithms, language models, finite state transducers, etc.,
- and enables harnessing aspects of the utterance that may be relevant for inferring the intent, but are not present in the text transcript, such as prosody.

End-to-end models have been made possible by deep learning, which automatically learns hierarchical representations of

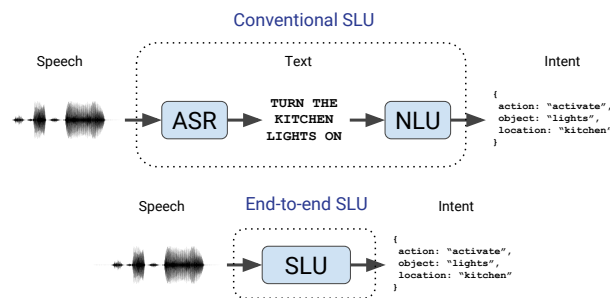


Figure 1: Conventional ASR → NLU system for SLU versus end-to-end SLU.

the input signal [10–14]. Speech is natural to represent in a hierarchical way: waveform → phonemes → morphemes → words → concepts → meaning. However, because speech signals are high-dimensional and highly variable even for a single speaker, training deep models and learning these hierarchical representations without a large amount of training data is difficult.

The computer vision [15, 16], natural language processing [17–21], and ASR [22–25] communities have attacked the problem of limited supervised training data with great success by pre-training deep models on related tasks for which there is more training data. Following their lead, we propose an efficient ASR-based pre-training methodology in this paper and show that it may be used to improve the performance of end-to-end SLU models, especially when the amount of training data is very small.

Our contributions are as follows:

- We introduce a dataset for realistic SLU experiments.
- We use this dataset to demonstrate effective speech model pre-training techniques for low-resource SLU.
- We make our code<sup>1</sup> and data<sup>2</sup> publicly available for researchers to replicate and build on our work.

## 2. Related work

Three key papers describing end-to-end SLU were written by Qian *et al.* [5], Serdyuk *et al.* [6], and Chen *et al.* [7]. Serdyuk *et al.* in [6] use no pre-training whatsoever. Qian *et al.* in [5] use an auto-encoder to initialize the SLU model. Chen *et al.* [7] pre-train the first stage of an SLU model to recognize graphemes; the softmax outputs of the first stage are then fed to a classifier second stage. The model proposed in this paper is similar to

<sup>1</sup>[github.com/lorenlugosch/pretrain\\_speech\\_model](https://github.com/lorenlugosch/pretrain_speech_model)

<sup>2</sup>[fluent.ai/research/fluent-speech-commands/](https://fluent.ai/research/fluent-speech-commands/)

theirs, but removes the restriction of the softmax bottleneck and uses alternative training targets, as we will describe later.

More recently, Haghani *et al.* in [8] compare four types of sequence-to-sequence models for SLU, including a direct model (end-to-end with no pre-training) and a multi-task model (uses a shared encoder whose output is ingested by a separate ASR decoder and SLU decoder). The model proposed here is somewhat similar to their multi-task model, although we do not use or require the ASR targets during SLU training.

The work listed above deals with very high resource SLU—in [8], for instance, the Google Home [26] dataset consists of 24 million labeled utterances. In contrast, Renkens *et al.* in [9] consider the problem of end-to-end SLU with limited training data, and find that capsule networks [27], compared to conventional neural network models, are more easily capable of learning end-to-end SLU from scratch. However, they do not consider the effect of pre-training on other speech data.

This previous work has all been conducted on datasets that are closed-source or too small to test hypotheses about the amount of data required to generalize well. The lack of a good open-source dataset for end-to-end SLU experiments makes it difficult for most people to perform high-quality, reproducible research on this topic. We therefore created a new SLU dataset, the “Fluent Speech Commands” dataset, which Fluent.ai has released along with this paper.

### 3. Dataset

This section describes the structure and creation of Fluent Speech Commands.

#### 3.1. Audio and labels

The dataset is composed of 16 kHz single-channel `.wav` audio files. Each audio file contains a recording of a single spoken English command that one might use for a smart home or virtual assistant, like “put on the music” or “turn up the heat in the kitchen”.

Each audio is labeled with three *slots*: action, object, and location. A slot takes on one of multiple *values*: for instance, the “location” slot can take on the values “none”, “kitchen”, “bedroom”, or “washroom”. We refer to the combination of slot values as the *intent* of the utterance. The dataset has 31 unique intents in total. We do not distinguish between domain, intent, and slot prediction, as is sometimes done in SLU [28].

The dataset can be used as a multi-label classification task, where the goal is to predict the action, object, and location labels. Since the slots are not actually independent of each other, a more careful approach would model the relationship between slots, e.g. using an autoregressive model, as in [8]. We use the simpler multi-label classification approach in this paper, so as to avoid the issues sometimes encountered training autoregressive models and instead focus on questions related to generalization using a simpler model. Alternately, the 31 distinct intents can be “flattened” and used as 31 distinct labels for a single-label classification task.

For each intent, there are multiple possible wordings: for example, the intent `{action: "activate", object: "lights", location: "none"}` can be expressed as “turn on the lights”, “switch the lights on”, “lights on”, etc.. These phrases were decided upon before data collection by asking employees at Fluent.ai, including both native and non-native English speakers, for various ways in which they might express a particular intent. There are 248 different phrases in total.

Table 1: *Information about the Fluent Speech Commands dataset.*

Split	# of speakers	# of utterances	# hours
Train	77	23,132	14.7
Valid	10	3,118	1.9
Test	10	3,793	2.4
Total	97	30,043	19.0

#### 3.2. Data collection

The data was collected using crowdsourcing. Each speaker was recorded saying each wording for each intent twice. The phrases to record were presented in a random order. Participants consented to data being released and provided demographic information about themselves. The demographic information about these anonymized speakers (age range, gender, speaking ability, etc.) is included along with the dataset.

The data was validated by a separate set of crowdsourcers. All audios deemed by the crowdsourcers to be unintelligible or contain the wrong phrase were removed. The total number of speakers, utterances, and hours of audio remaining is shown in Table 1.

#### 3.3. Dataset splits

The utterances are randomly divided into train, valid, and test splits in such a way that no speaker appears in more than one split. Each split contains all possible wordings for each intent, though our code has the option to include data for only certain wordings for different sets, to test the model’s ability to recognize wordings not heard during training. The dataset has a `.csv` file for each split that lists the speaker ID, file path, transcription, and slots for all the `.wav` files in that split.

#### 3.4. Related datasets

Here we review some related public datasets and show the gap that Fluent Speech Commands fills.

The Google Speech Commands dataset [29] (to which the name “Fluent Speech Commands” is an homage) is a free dataset of 30 single-word spoken commands (“yes”, “no”, “stop”, “go”, etc.). This dataset is suitable for keyword spotting experiments, but not for SLU.

ATIS is an SLU dataset consisting of utterances related to travel planning<sup>3</sup>. This dataset can only be obtained expensively from the Linguistic Data Consortium.

The Grabo, Domotica, and Patcor datasets are three related datasets of spoken commands for robot control and card games developed by KU Leuven<sup>4</sup> and used in [9]. These datasets are free, but have only a small number of speakers and phrases.

The free audio-based Snips SLU Dataset [30] is quite close to what we want. This dataset has a great variety of phrases, but less audio—2,946 English utterances and 1,138 French utterances, compared to 30,043 in Fluent Speech Commands.

In contrast to these datasets, Fluent Speech Commands is simultaneously audio-based, reasonably large, and free, and contains several multiple-word commands corresponding to each of the intents, with multiple recordings for each wording.

<sup>3</sup><https://catalog.ldc.upenn.edu/LDC94S19>

<sup>4</sup><https://www.esat.kuleuven.be/psi/spraak/downloads/>

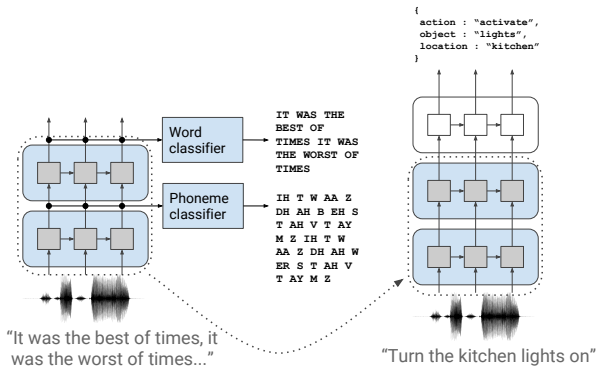


Figure 2: The lower layers of the model are pre-trained using ASR targets (words and phonemes). The word and phoneme classifiers are discarded, and the features from the pre-trained part of the model (blue) are used as the input to the subsequent module (white), which is trained using SLU targets.

## 4. Model and Pre-training Strategy

Using the dataset described in the previous section, we test the performance of our proposed model and pre-training strategy (Fig. 2), which we describe here. The model is a deep neural network consisting of a stack of modules, where the first modules are pre-trained to predict phonemes and words. The word and phoneme classifiers are discarded, and the entire model is then trained end-to-end on the supervised SLU task. Here we justify and give more details about these design decisions.

### 4.1. Which ASR targets to use?

ASR models are trained using a variety of targets, including phonemes, graphemes, wordpieces, or more recently whole words [31–33]. We choose whole words as the pre-training targets, since this is what a typical NLU module would expect as input. A large vocabulary ASR dataset contains too many unique words (LibriSpeech [34] has more than 200,000) to efficiently assign an output to each one; we only assign a label to the 10,000 most common words. This leaves much of the pre-training data without any labels, which wastes data. By using phonemes as intermediate pre-training targets [20, 35, 36], we are able to pre-train on speech segments with no word label. Additionally, we find that using phonemes as intermediate targets speeds up word-level pre-training [37–39].

We use the Montreal Forced Aligner [40] to obtain word- and phoneme-level alignments for LibriSpeech, and we pre-train the model on the entire 960 hours of training data using these alignments<sup>5,6</sup>. Using force-aligned labels has the additional benefit of enabling pre-training using short, random crops rather than entire utterances, which reduces the computation and memory required to pre-train the model.

### 4.2. Phoneme module

The first module takes as input the audio signal and outputs  $\mathbf{h}^{\text{phoneme}}$ , a sequence of hidden representations that are pre-trained to predict phonemes. The phoneme-level logits are com-

<sup>5</sup>Our alignments can be downloaded from <https://zenodo.org/record/2619474#.XKDP2VnKg1g>

<sup>6</sup>We use textgrid.py to process the alignments. <https://github.com/kylebgorman/textgrid>

puted using a linear classifier:

$$\mathbf{l}^{\text{phoneme}} = W^{\text{phoneme}} \mathbf{h}^{\text{phoneme}} + b^{\text{phoneme}}. \quad (1)$$

The phoneme module is implemented using a SincNet layer [41, 42], which processes the raw input waveform, followed by multiple convolutional layers and recurrent layers with pooling (to reduce the sequence length) and dropout. More detailed hyperparameters can be found in our code.

### 4.3. Word module

The second module takes as input  $\mathbf{h}^{\text{phoneme}}$  and outputs  $\mathbf{h}^{\text{word}}$ . Similar to the phoneme-level module, it uses recurrent layers with dropout and pooling, and is pre-trained to predict whole word targets using another linear classifier:

$$\mathbf{l}^{\text{word}} = W^{\text{word}} \mathbf{h}^{\text{word}} + b^{\text{word}}. \quad (2)$$

Notice that the input to this module is  $\mathbf{h}^{\text{phoneme}}$ , not  $\mathbf{l}^{\text{phoneme}}$ , and likewise the output to the next stage is  $\mathbf{h}^{\text{word}}$ , not  $\mathbf{l}^{\text{word}}$ . There are two good reasons for forwarding  $\mathbf{h}$  instead of  $\mathbf{l}$ . The first is that we don’t want to remove a degree of freedom from the model: the size of  $\mathbf{l}$  is fixed by the number of targets, and this would in turn fix the size of the next layer of the model. The second reason is that computing  $\mathbf{l}^{\text{word}}$  requires multiplying and storing a large ( $\approx 2.5$  million parameters) weight matrix, and by discarding this matrix after pre-training, we save on memory and computation.

### 4.4. Intent module

The third module, which is not pre-trained, maps  $\mathbf{h}^{\text{word}}$  to the predicted intent. Depending on the structure of the intent representation, the intent module might take on various forms. Since in this work we use a fixed three-slot intent representation, we implement this module using a recurrent layer, followed by max-pooling to squash the sequence of outputs from the recurrent layer into a single vector of logits corresponding to the different slot values, similar to [6].

### 4.5. Unfreezing schedule

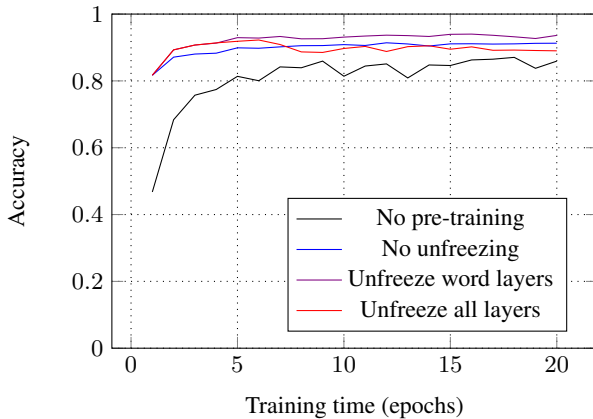
Although the pre-trained model works well as a frozen feature extractor, it may be preferable to “unfreeze” its weights and finetune them for the SLU task with backpropagation. Similar to ULMFiT [18], we find that gradually unfreezing the pre-trained layers works better than unfreezing them all at once. We unfreeze one layer each epoch, and stop at a pre-determined layer, which is a hyperparameter.

## 5. Experiments

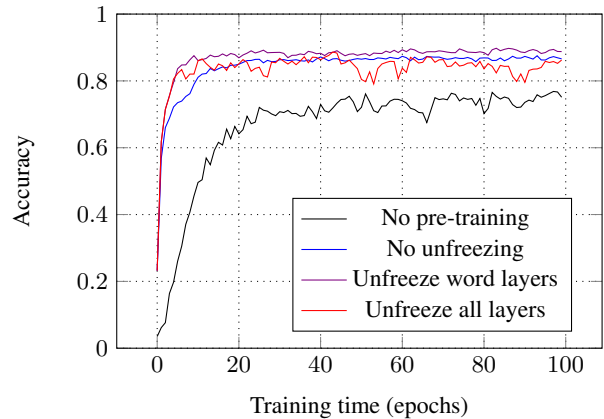
Here we report results for three experiments on Fluent Speech Commands: using the full dataset, using a subset of the dataset, and using a subset of wordings.

### 5.1. Full dataset

We first trained models given the entire SLU training set. The models used one of: 1) no pre-training (randomly initialized), 2) pre-training with no unfreezing, 3) gradually unfreezing only the word layers, or 4) gradually unfreezing both the word layers and phoneme layers. What we report here as “accuracy” refers to the accuracy of all slots for an utterance taken together—that is, if the predicted intent differs from the true intent in even one slot, the prediction is deemed incorrect.



(a) Full dataset.



(b) 10% of the dataset.

Figure 3: Accuracy on the validation set over time for models trained on (a) the full SLU dataset or (b) 10% of the dataset.

Table 2: Accuracy on the test set for different models, given the full training dataset or a 10% subset of the training data.

Model	Accuracy (full)	Accuracy (10%)
No pre-training	96.6%	88.9%
No unfreezing	98.8%	97.9%
Unfreeze word layers	98.7%	97.9%
Unfreeze all layers	97.2%	95.8%

The validation accuracy of these models over time is shown in Fig. 3a. The best results are obtained when only the word layers of the pre-trained model are unfrozen. This may be because the fully unfrozen model begins to forget the more general phonetic knowledge acquired during pre-training. For the test set, the frozen model and partially unfrozen model perform roughly equally well (Table 2, “full” column), possibly because the test set is “easier” than the validation set. In all cases, the pre-trained models outperform the randomly initialized model.

## 5.2. Partial dataset

To simulate a smaller dataset, we randomly selected 10% of the training set, and used this instead of the entire training set. Fig. 3b shows the validation accuracy (on the entire validation set, not a subset) over time. A similar trend is observed as for the entire dataset: unfreezing the word layers works best. The gap in final test accuracy between the randomly initialized model and the pre-trained models increases (Table 2, “10%” column); the final test accuracy for the pre-trained models drops only slightly, further highlighting the advantage of our proposed method.

## 5.3. Generalizing to new wordings

What happens if new wordings appear in the test data that never appear in the training data? This is an important question, since it is generally impractical to try to imagine every possible wording for a particular intent while gathering training data.

To test this, we trained models on three specific phrases, “turn on the lights”, “turn off the lights”, and “switch on the lights” (273 utterances total), and tested on those same phrases, as well as a new phrase: “switch off the lights”. If the model incorrectly infers that utterances that contain “switch” always

correspond to turning on the lights, it will incorrectly guess that “switch off the lights” corresponds to turning on the lights; if the model infers that the presence of the word “off” corresponds to turning off the lights, it will generalize to the new phrase. The randomly initialized model was unable to fit this tiny training set, even with a very low learning rate and no regularization. The pre-trained models were able to generalize to the new wording (with 97% accuracy on the validation set, which contains more examples of the new phrase than of the training phrases).

However, there are many situations in which our model does not correctly generalize. For example, if the model is trained only with examples containing “bedroom” and “washroom”, but then tested on an example containing “bathroom”, it will guess the intent corresponding to “bedroom” because “bedroom” sounds more similar to “bathroom” than to “washroom”, even though “washroom” is the correct meaning. In text-based NLU, this scenario can be handled using word embeddings, which represent words in such a way that words with similar meanings have similar vector representations [2, 43]. It may be possible to teach the pre-trained part of the model to output “embedding-like” word representations so that the intent module can recognize the meaning of phrases with synonyms.

## 6. Conclusion

In this paper, we proposed a pre-training methodology for end-to-end SLU models, introduced the Fluent Speech Commands dataset, and used this dataset to show that our pre-training techniques improve performance both for large and small SLU training sets. In the future, we plan to continue using Fluent Speech Commands to explore the limitations of end-to-end SLU, like new wordings and synonyms not observed in the SLU dataset, to see if these limitations can be overcome.

## 7. Acknowledgements

We would like to acknowledge the following for research funding and computing support: NSERC, Calcul Québec, Compute Canada, the Canada Research Chairs, and CIFAR.

Thanks to Dima Serdyuk and Kyle Kastner at Mila, and Farzaneh Fard, Luis Rodriguez Ruiz, Sam Myer, Mohamed Mhiri, and Arash Rad at Fluent.ai for helpful discussions with us about this work.

## 8. References

- [1] G. Tur and R. D. Mori, *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*. Wiley, 2011.
- [2] G. Mesnil, Y. Dauphin, K. Yao, Y. Bengio, L. Deng, D. Hakkani-tur, and X. He, "Using Recurrent Neural Networks for Slot Filling in Spoken Language Understanding," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2015.
- [3] A. Coucke, A. Saade, A. Ball, T. Bluche, A. Caulier, D. Leroy, C. Doumouro, T. Gisselbrecht, F. Caltagirone, T. Lavril, M. Primet, and J. Dureau, "Snips Voice Platform: an embedded Spoken Language Understanding system for private-by-design voice interfaces," *arXiv:1805.10190v3*, 2018.
- [4] A. L. Gorin, G. Riccardi, and J. H. Wright, "How may I help you?" *Speech Communication*, vol. 23, no. 1-2, pp. 113-127, 1997.
- [5] Y. Qian, R. Ubale, V. Ramanarayanan, and P. Lange, "Exploring ASR-free end-to-end modeling to improve spoken language understanding in a cloud-based dialog system," *ASRU*, 2017.
- [6] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, "Towards end-to-end spoken language understanding," *ICASSP*, 2018.
- [7] Y.-P. Chen, R. Price, and S. Bangalore, "Spoken language understanding without speech recognition," *ICASSP*, 2018.
- [8] P. Haghani, A. Narayanan, M. Bacchiani, G. Chuang, N. Gaur, P. Moreno, R. Prabhavalkar, Z. Qu, and A. Waters, "From Audio to Semantics: Approaches to end-to-end spoken language understanding," *arXiv:1809.09190*, 2018.
- [9] V. Renkens and H. Van hamme, "Capsule Networks for Low Resource Spoken Language Understanding," *Interspeech*, 2018.
- [10] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [11] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278-2324, 1998.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *NeurIPS*, 2012.
- [13] A. Graves and N. Jaitly, "Towards End-to-End Speech Recognition with Recurrent Neural Networks," *ICML*, 2014.
- [14] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanowski, A. Coates, G. Diamos *et al.*, "Deep Speech 2: End-to-End Speech Recognition in English and Mandarin," *ICML*, 2015.
- [15] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" *NeurIPS*, 2014.
- [16] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *ICLR*, 2015.
- [17] A. M. Dai and Q. V. Le, "Semi-supervised sequence learning," *NeurIPS*, 2015.
- [18] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," *ACL*, 2018.
- [19] A. Radford and T. Salimans, "Improving Language Understanding by Generative Pre-Training," *Technical report, OpenAI*, 2018.
- [20] V. Sanh, T. Wolf, and S. Ruder, "A hierarchical multi-task approach for learning embeddings from semantic tasks," *AAAI*, 2019.
- [21] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional Transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [22] D. Wang and T. Zheng, "Transfer learning for speech and language processing," *APSIPA*, 2015.
- [23] G. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large vocabulary speech recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 1, pp. 30-42, 2012.
- [24] J. Kunze, L. Kirsch, I. Kurenkov, A. Krug, J. Johannsmeier, and S. Stober, "Transfer Learning for Speech Recognition on a Budget," *ACL*, 2017.
- [25] P. Ghahremani, V. Manohar, H. Hadian, D. Povey, and S. Khudanpur, "Investigation of transfer learning for ASR using LF-MMI trained neural networks," *ASRU*, 2017.
- [26] B. Li, T. N. Sainath, A. Narayanan, J. Caroselli, M. Bacchiani, A. Misra, I. Shafran, H. Sak, G. Pundak, K. K. Chin *et al.*, "Acoustic Modeling for Google Home," *Interspeech*, 2017.
- [27] S. Sabour, N. Frosst, and G. E. Hinton, "Dynamic routing between capsules," *NeurIPS*, 2017.
- [28] Y.-Y. Wang, L. Deng, and A. Acero, "Spoken language understanding," *IEEE Signal Processing Magazine*, vol. 22, no. 5, pp. 16-31, 2005.
- [29] P. Warden, "Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition," *arXiv:1804.03209v1*, 2018.
- [30] A. Saade, A. Coucke, A. Caulier, J. Dureau, A. Ball, T. Bluche, D. Leroy, C. Doumouro, T. Gisselbrecht, F. Caltagirone, T. Lavril, and M. Primet, "Spoken Language Understanding on the Edge," *arXiv:1805.10190v3*, 2018.
- [31] H. Soltau, H. Liao, and H. Sak, "Neural Speech Recognizer: Acoustic-to-Word LSTM Model for Large Vocabulary Speech Recognition," *arxiv:1610.09975*, 2016.
- [32] K. Audhkhasi, B. Ramabhadran, G. Saon, M. Picheny, and D. Nahamoo, "Direct Acoustics-to-Word Models for English Conversational Speech Recognition," *Interspeech*, 2017.
- [33] K. Audhkhasi, B. Kingsbury, B. Ramabhadran, G. Saon, and M. Picheny, "Building competitive direct acoustics-to-word models for english conversational speech recognition," *ICASSP*, 2018.
- [34] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," *ICASSP*, 2015.
- [35] S. Fernández, A. Graves, and J. Schmidhuber, "Sequence Labelling in Structured Domains with Hierarchical Recurrent Neural Networks," *IJCAI*, 2007.
- [36] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, no. 1, pp. 41-75, 1997.
- [37] T.-S. Nguyen, S. Stueker, and A. Waibel, "Using multi-task learning to improve the performance of acoustic-to-word and conventional hybrid models," *arXiv:1902.01951*, 2019.
- [38] R. Sanabria and F. Metze, "Hierarchical multitask learning with CTC," *SLT*, 2018.
- [39] K. Krishna, S. Toshiwal, and K. Livescu, "Hierarchical Multitask Learning for CTC-based Speech Recognition," *arXiv:1807.06234*, 2018.
- [40] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal Forced Aligner: Trainable text-speech alignment using Kaldi," *Interspeech*, 2017.
- [41] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with SincNet," *SLT*, 2018.
- [42] —, "Interpretable convolutional filters with SincNet," *NeurIPS Workshop on Interpretability and Robustness for Audio, Speech and Language (IRASL)*, 2018.
- [43] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.