



Large scale weakly and semi-supervised learning for low-resource video ASR

Kritika Singh*, Vimal Manohar*, Alex Xiao*, Sergey Edunov, Ross Girshick, Vitaliy Liptchinsky, Christian Fuegen, Yatharth Saraf, Geoffrey Zweig, Abdelrahman Mohamed

Facebook AI

{skritika, vimalmanohar, axiao, edunov, rbg, vitaliy888, fuegen, ysaraf, gzweig, abdo}@fb.com

Abstract

Many semi- and weakly-supervised approaches have been investigated for overcoming the labeling cost of building high-quality speech recognition systems. On the challenging task of transcribing social media videos in low-resource conditions, we conduct a large scale systematic comparison between two self-labeling methods on one hand, and weakly-supervised pre-training using contextual metadata on the other. We investigate distillation methods at the frame level and the sequence level for hybrid, encoder-only Connectionist Temporal Classification (CTC) based, and encoder-decoder speech recognition systems on Dutch and Romanian languages using 27,000 and 58,000 hours of unlabeled audio respectively. Although all approaches improved upon their respective baseline word error rates (WER) by more than 8%, sequence-level distillation for encoder-decoder models provided the largest relative WER reduction of 20% compared to the strongest data-augmented supervised baseline.

Index Terms: Weak and self-labeling, speech recognition

1. Introduction

Recent advances in speech recognition systems have enabled successful large scale deployments of various customer-facing speech applications, e.g. personal conversational agents, automatic transcriptions for accessibility, and multi-modal video understanding. This trend has increased the need for developing accurate automatic speech recognition (ASR) models for many languages and domains. However, there are significant challenges to achieving this.

First, even though current ASR systems are arguably within striking distance of human performance for broadcast news and telephone-speech domains [1, 2, 3], more challenging real-world scenarios involving unconstrained, natural speech that is filled with background music and noise, various speaking styles and emotions, disfluencies, heavy accents, un-cued speaker and language switching, is still an open problem for speech recognition systems [4]. In this paper, we focus on the domain of public social media videos which involve all these challenges while representing an interesting benchmark for evaluating the effectiveness of different learning methods due to their ever increasing amount, multi-modal nature, and the availability of related metadata, e.g. video title, post text, and comments.

Second, it can be prohibitively expensive and difficult to collect sufficient amounts of supervised training labels to feed data hungry neural speech recognition models for each language. Therefore, we'd like to minimize the amount of supervised data used in training these models. With increased access to large computational resources, three families of methods have emerged in this direction: (1) Using large volumes of

unpaired audio and text data. [5, 6, 7, 8] (2) Augmenting audio data with contextual metadata as distant labels [9, 10, 11, 12]. (3) Data augmentation through reverberation, structured noise, speed perturbation, time and frequency masking [13, 14]. Given the various existing approaches for combining unlabeled audio with unpaired text and for incorporating distant labels, it is unclear, however, which of these methods are complementary, how they compare to each other, and how they can be combined effectively into a scalable recipe to maximize speech recognition performance for a complex domain like social media videos under low-resource constraints.

In this paper, we take a step towards answering some of these questions by conducting large scale experiments to compare and combine methods drawn from these families of techniques, focusing on low-resource setups for transcribing public social media videos in two languages: Dutch and Romanian, using 27,000 and 58,000 hours of unlabeled data respectively. Applying data augmentation for all models [13, 14], we compare a recently proposed weakly-supervised approach [12] (Section 2) and two commonly used self-labeling methods: frame-level distillation [15] and sequence-level distillation [5, 16] for hybrid, CTC-based, and encoder-decoder ASR setups (Section 3).

2. Weakly-supervised Speech Recognition

Following [12], we use video metadata, e.g. title and post text, as distant labels for the ASR task. We use two sets of training data: (i) $\{X, Y^s\} \in \mathcal{D}^s$ is the supervised set where X and Y^s are pairs of audio features and label sequences. (ii) $\{X, Y^w\} \in \mathcal{D}^w$ is the weakly-supervised set where X and Y^w are pairs of audio features and the corresponding contextual text. The targets Y^s and Y^w are sequences of sub-word units [17].

An encoder-decoder approach [18, 19] is used for maximizing the conditional probability of generating the contextual text sequence Y^w given X , an input sequence of mel-scale log filterbank features, where $x_i \in R^d$

$$\begin{aligned} \mathcal{F}^w &= p(Y^w | X; \theta^w) \\ &= \prod_{i=1}^M p(y_i^w | y_1^w, y_2^w, \dots, y_{i-1}^w, x_1, x_2, \dots, x_T; \theta^w) \end{aligned} \quad (1)$$

The attention-based encoder-decoder framework fits well with the weak supervision since it offers flexible alignment and unconstrained coverage between input and output sequences. Other ASR training approaches aren't as suitable given the abstractive relationship between Y^w and X . Hybrid Hidden Markov Model-Neural Network (HMM-NN) and CTC approaches assume either a frame-level or a monotonic alignment between input and target sequences, and they constrain the possible length of the target sequence by the input sequence length.

The final objective function is $\mathcal{F} = \mathcal{F}^w + \mathcal{F}^s$ where the supervised term, $\mathcal{F}^s = p(Y^s | X; \theta^s)$, is also maximized using

*Equal contribution

the encoder-decoder approach. We share the full model for both types of data, where $\theta^w = \theta^s = \{\theta_{enc}, \theta_{dec}\}$ combines the parameters in the audio encoder and the language generation and attention modules in the decoder. During training, we alternate, with some mixing ratio, between mini-batches sampled from the two training sets \mathcal{D}^s and \mathcal{D}^w .

3. Self-labeled Speech Recognition

Self-labeling is one of the most effective methods of semi-supervised learning for speech recognition [20, 9, 5], where a teacher model with limited supervision extends the training data by labeling an extra amount of unlabeled data. In this study, we investigate two approaches:

3.1. Frame-level distillation

In hybrid ASR modeling, acoustic models are trained to map input audio features into phonetic or graphemic output units by minimizing the frame-level cross entropy between the model predictions and ground truth labels. Frame-level knowledge distillation [21, 22] replaces the ground truth labels with probability distributions generated from a teacher model over all output units for the same set of supervised data, or for a large volume of unlabeled audio [15]. It minimizes the KL divergence between outputs of the model, referred to as a student, and the teacher model’s frame-level distributions:

$$\mathcal{F}^d = -\frac{1}{T} \sum_{t=1}^T \sum_{k=1}^K p(y_t = k|X) \log \hat{p}(y_t = k|X), \quad (2)$$

where X is the sequence of features with T time-steps, K is the number of classes, $p(y_t = k|X)$ is the teacher frame-level probability for class k and time step t , and $\hat{p}(y_t = k|X)$ is the student frame-level probability.

3.2. Sequence-level distillation

In sequence-level distillation, the student model attempts to learn the teacher model’s distribution over the entire sequence as opposed to just the individual frames. The full teacher sequence distribution is often intractable, but it can be approximated with a lattice [23] for hybrid or the top k hypotheses for CTC [24] and seq2seq [25, 26]. Here, we consider an approximation by just training on the top-1 hypotheses generated by the teacher model. Unlike frame-level distillation, sequence-level distillation utilizes a language model (LM) during label generation. The LM can be adapted to a specific genre or domain, and can also be used to filter out audio segments with very unlikely word sequences. Sequence-level distillation is the first form of self-labelling applied for speech recognition [20, 9, 27, 5]. It’s also commonly used for non-autoregressive machine translation systems [28].

When using target word sequences instead of frame-level probability distributions, the student system has greater flexibility to (a) Follow any modelling paradigm, e.g. hybrid, encoder-only with CTC loss, or attention/transducer-based encoder-decoder (b) Use an output vocabulary for the student that is different from the teacher ASR system, (c) Filter the unlabeled data and focus on specific weaknesses in the originally collected labeled data, e.g. improving recognition of named entities in a specific domain [16] and, (d) Evaluate how synergistic self-labels and weak-supervision in the form of metadata labels are through multi-task learning of the encoder-decoder model with mini-batches sampled from both data sources which we will show in the experimental section.

4. Experiments

4.1. Data

We explore the aforementioned weakly-supervised and self-labeling approaches on two languages under low-resource conditions: Dutch and Romanian. Both languages have three test sets – test-clean, test-noisy and test-extreme – listed in increasing order of difficulty in acoustic conditions, and a single dev-noisy set. The sizes of the supervised sets – train, dev-noisy, test-clean, test-noisy and test-extreme – are 290, 13, 10, 25, and 12 hours respectively for Dutch. They are 163, 6, 6, 12, and 10 hours respectively for Romanian. All our datasets are sampled from public videos that are anonymized. The unlabeled audio data for both languages come from videos that are up to 5 minutes in duration. First, the whole data is decoded with a baseline hybrid model and segmented using hypothesized word boundaries into chunks at most 10s in duration while removing non-speech portions.

This process yields the `data-large` set, which is about 27,000 hours for Dutch and 58,000 hours for Romanian. For our meta-data supervision experiments, these sets are further filtered to preserve videos with meta-data between 50 and 700 characters. Videos with meta-data that are completely orthogonal to the baseline ASR hypotheses (zero word overlap) are discarded. This process produces our weak-supervision data set `data-ws` of 7,900 hours for Dutch and 6,400 hours for Romanian. Note that the full metadata text is shared for different segments of the same video during weakly-supervised learning.

4.2. Experimental setup

4.2.1. Data Preprocessing

We use 80 dimensional mel-scale log filterbank features computed every 10ms over 25ms windows for hybrid models and 16ms for all other models. Both labeled and unlabeled data are augmented with copies using 0.9 and 1.1 speed perturbation [13], and then one more copy by superposition with 10s clips of diverse noise events. In all training stages, we apply on-the-fly time and frequency masking following the Switchboard Mild (SM) policy from [14]¹.

4.2.2. System Description

To facilitate comparisons, all our hybrid, CTC-based, and encoder-decoder models have the same encoder architecture of 5 layer, 800 hidden units latency controlled bidirectional LSTM (LC-BLSTM) with a factor of 2 sub-sampling over time after the first layer. While an architecture search and/or using a different neural model, e.g. Transformers [29, 30, 31, 32], may lead to better absolute WER numbers, we believe that the conclusions reached in our experiments are invariant to such modeling decisions.

The hybrid systems use 8800 output cheneone units [33] for Dutch and 9272 for Romanian models. The CTC-based and encoder-decoder models use 1000 sub-word output units [17] trained on the supervised transcripts.

For decoding, both the hybrid and CTC-based models use a Weighted Finite State Transducer (WFST) decoder with a 5-gram LM trained on the supervised transcripts. Following the architecture in [12], the decoder module of the encoder-decoder models consists of 4 1-D convolutional layers with kernel size

¹Our self-labeled hybrid models don’t use speed and noise perturbation as initial experiments found them of limited benefit given the unlabeled data size.

of 3 and 256 output features followed by 2 transformer blocks with 1k hidden dimensions, 16 self-attention heads and a 4k projection layer before the ReLU nonlinearity. At inference time, a beam search with a beam size of 20 is used, with no language model fusion, to get the 1-best hypothesis.

4.2.3. Training Details

The encoder-decoder based self-labeling and weakly-supervised training experiments consist of three phases as described in [12]: (1) An initial supervised `burn-in` phase (2) A `train-main` phase that uses a mixture of supervised and the unlabeled data with either meta-data or self-labels as targets. (3) A final supervised-only `fine-tune` phase. During `train-main` we save checkpoints every 5k model updates and average the last 20 checkpoints to initialize the `fine-tune` phase. We use the same process for self-labeling using CTC. For both encoder-decoder and CTC, we use Adam [34] optimizer with a fixed learning rate of 4×10^{-4} during `burn-in` and `train-main` phases, and 4×10^{-5} during encoder-decoder `fine-tune` and 5×10^{-5} during CTC `fine-tune` phase. Gradient norm clipping at 10.0 is used, where total gradients are scaled by the number of utterances in each mini-batch.

Our hybrid acoustic models are first trained with frame-level cross entropy (CE) then with the Lattice-free Maximum Mutual Information (LFMMI) [35] criterion where frame-level supervision and numerator lattices are generated using a smaller bootstrapping hybrid system. We use a dropout of 0.5 and Adam optimizer with learning rates of 0.005 and 0.00001 for the CE and LFMMI stages with Incremental block distributed data parallelism (BMUF) [36]. The learning rate is halved when the held-out set doesn't improve at the end of each epoch. The supervised baseline is used for segmenting the unlabeled data and generating the self-labels for sequence-distillation experiments.

We chose the size of the teacher model for frame-level distillation experiments based on performance: 6 layers of LC-BLSTMs with 1000 hidden units for Dutch and 5 layers of 800 hidden units for Romanian. Following [15], we approximate the full teacher distribution with the top 3 predictions per frame, to minimize storage and network costs, which on average was enough to cover 99% of the per-frame probability mass. After distillation, hybrid models are then fine-tuned with LFMMI on supervised data.

4.3. Results

Table 1 compares the two strategies for self-labeling on `data-large`, frame-level distillation for hybrid models and sequence-level distillation for all of hybrid, CTC-based and encoder-decoder models, with respect to their supervised baselines. Our strongest baselines for both Dutch and Romanian are the LFMMI fine-tuned hybrid models. For sequence-level distillations, all self-labeled systems provided relative improvements over their respective baselines with more than 12% for hybrid, 20% for CTC-based, and 37% for encoder-decoder. The relative improvements on the Romanian language are on average 8% higher than Dutch, which has about double the supervised data size, across all approaches.

The encoder-decoder model has two advantages over the other two systems. First, by relying entirely on sub-word units, it is an open vocabulary system with the ability to generate out-of-vocabulary words. Moreover, it learns a longer range lan-

	Dutch			Romanian		
	clean	noisy	extreme	clean	noisy	extreme
Supervised baseline						
Hybrid (LFMMI)	23.6	23.3	32.8	17.5	19.5	32.9
CTC	26.7	26.3	36.8	20.8	22.2	38.3
Enc-Dec	27.2	27.0	39.0	25.5	26.8	46.0
Self-labeling using frame-level distillation						
Top 3 (CE)	23.3	22.8	32.3	15.7	17.9	29.6
+ LFMMI	21.3	20.7	29.8	14.7	17.0	28.8
Top 1 (CE)	22.9	22.6	32.0	15.9	17.8	29.7
+ LFMMI	21.7	21.6	30.9	14.7	16.9	28.6
Self-labeling using sequence-level distillation						
Hybrid (CE)	23.6	23.3	32.6	16.3	18.4	30.7
+ LFMMI	20.9	20.8	29.7	14.7	16.8	28.3
SL-LFMMI	22.1	21.8	31.4	15.6	17.6	29.5
CTC	22.5	22.2	31.4	14.9	17.2	29.4
Enc-Dec	18.4	18.5	27.9	13.1	15.6	27.3

Table 1: WERs of different self-labeling approaches and ASR models on `data-large` relative to baseline supervised models on the three test sets for Dutch and Romanian. All + LFMMI models were seeded with the CE model from the previous row and use supervised data only except for the LFMMI row prefixed with SL- that uses self-labeling data.

guage model on these sub-word units over the self-labels while observing the entire encoder output. Therefore, it provides 20% on average compared the best baseline supervised system.

Within sequence-level distillation systems, self-labeled CTC-trained models are better than the self-labeled CE-trained hybrid model by 5% on average but with LFMMI discriminative training the hybrid system leaps forward by about 4%. Adding discriminative training techniques like sMBR [37] on top of the CTC-trained model would bring further gains which we leave for a future investigation. Confirming its sensitivity to transcription quality [38, 15], using self-labels for LFMMI training, even with the large amount of audio data, is not as useful as using the much smaller supervised data in the discriminative training stage of hybrid models.

Interestingly, using Top-3 frame-level labels for distillation provided the same level of gains as the hybrid sequence-level distillation which utilizes a language model during label generation. When we use only the Top-1, the performance stays almost the same for Romanian but degrades slightly after LFMMI fine-tuning in Dutch.

For the encoder-decoder models, table 2 compares weakly supervised learning using `data-ws` with meta-data targets and self-labeling using both `data-ws` and `data-large` audio. In both cases, a supervised `fine-tune` stage is used at the end of training. While weakly-supervised models improves over their respective baseline models from table 1 by about 10%, they don't come near observed gains from self-labeling. For clean and noisy conditions the smaller audio size of `data-ws` is enough to get all the gains from self-labeling. However, performance in extremely noisy conditions continues to improve with larger volume of unlabeled audio for both languages. Combining weak-supervision on `data-ws` and self-labeling on `data-large` barely makes any noticeable change compared to the self-labeling alone, showing that the model has achieved all its gains from diverse audio conditions without uti-

	Dutch			Romanian		
	clean	noisy	extreme	clean	noisy	extreme
Weak-supervision (WS)	24.5	24.6	35.9	21.6	25.9	42.5
SL (data-ws)	18.4	18.8	28.7	13.3	15.6	28.2
SL (data-large)	18.4	18.5	27.9	13.1	15.6	27.3
0.7 SL (data-large)						
+ 0.3 WS	18.5	18.6	28.1	13.1	15.4	28.0

Table 2: WERs of weakly-supervised (WS) and self-labeled (SL) enc-dec models using both data-ws and data-large as well as their combination.

	Dutch		
	clean	noisy	extreme
0-iteration 80M parameters	23.6	23.3	32.8
0-iteration 150M parameters	23.2	23.1	33
0-iteration 240M parameters	23.5	23.2	32.9
1-iteration 80M parameters	21.3	20.7	29.8
2-iteration 150M parameters	20.5	20.0	28.6
3-iteration 240M parameters	19.7	19.4	27.6

Table 3: WERs of the Dutch test sets using Top-3 iterative frame-level distillation for up to three iterations of relabeling data-large.

lizing much of the information in the textual meta-data. This can be attributed to the meta-data quality and their weak relevance to audio content.

Motivated by the success of early iterative self-labeling [5] and more recently in the computer vision tasks [39], table 3 demonstrates iteratively distilling progressively larger student models that are distilled as teachers in following steps. Each iteration accounts for a round of frame-level distillation of a hybrid model followed by LFMFI fine-tuning using supervised data. In 0-iteration cases, baseline models of different sizes are trained on the data-augmented supervised data with almost no difference in WER. Each iteration of self-labeling added about 3% to 4% reduction in WER compared to a single iteration, pushing the overall gain from 10% to 16% on average compared to the baseline supervised system. Larger student models have greater capacity to benefit from the diverse audio conditions in the data-large set.

5. Discussion and Related Work

Reducing the amount of manual transcription for building speech recognition systems has been a constant research theme over the past three decades [40, 20, 9, 5, 41, 42] given the revolutionary increase and access to unlabeled data and computational resources. Starting as far as the mid-1990s, early research work used a bootstrap model trained on few hours of supervised data to generate labels for a larger set of audio data, with confidence filters applied to remove wrong transcriptions [40, 20, 43, 27]. TV shows and news closed captions were one of the early forms of audio metadata used either to bias the LM during decoding or as direct labels [44, 9, 45, 46]. Visual grounding used as weak-labels for learning audio representations in [10, 11]. Starting from as little as 1.2 hours of transcribed data, [5] used iterative re-labeling and self-labeling

while progressively increasing the model size to halve the overall system WER. Motivated by the work on model compression [21] and distillation [47, 22], and using 1 million of unlabeled personal assistant data, [15] showed relative WER reduction of about 10% to 20% compared to supervised baseline. Incorporating the LM information within a sequence discriminative loss function for self-labeling was proven to be a difficult task given the noise in labels [46, 48, 38]. On the other hand, seq2seq models showed greater grace working with noisy labels both for personal assistant domains [16] and audio books [49, 50]. With the full 960h Librispeech [51] data for building a strong initial model, [52] showed that using the larger Libri-light [53] unlabeled data in the same domain in a self-labeling seq2seq setup significantly improves over the best supervised system achieving a new state-of-the-art for Librispeech.

6. Conclusion and Future work

In this paper, we investigated self- and weakly-supervised training with hybrid, encoder-decoder and CTC-based models for low-resource speech recognition on Dutch and Romanian public social media videos with about 300 and 150 hours of supervised data respectively. Using 27,000 hours and 58,000 hours of unlabeled Dutch and Romanian data, self-labeled encoder-decoder speech recognition models with sequence-level distillation achieved WER reduction of more than 20% relative to the strongest data-augmented baseline. While weakly-supervised trained models using video meta-data as distant labels brought 10% relative improvement over the encoder-decoder baseline, they didn't come close to the self-labeling gains. Combining the distant and self-labels barely made any difference, showing that the combined system didn't utilize the contextual meta-data information. We additionally observed that on hybrid models, frame-level distillation using Top-3 frame-level labels provided the same level of gains as the sequence-level distillation which utilizes a language model during label generation. While the encoder-decoder models seem to benefit from the implicit sub-word language model of the self-labels, further research is needed to realize these benefits on other systems. Moving beyond our current minimum of 150 hours of supervised data, future work will focus on methods that preserve the same level of WER while using one or two orders of magnitude lower volumes of supervised data.

7. References

- [1] 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 *ICML*, 2016.
- [2] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, B. Ramabhadran, M. Picheny, L.-L. Lim *et al.*, "English conversational telephone speech recognition by humans and machines," *arXiv preprint arXiv:1703.02136*, 2017.
- [3] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Achieving human parity in conversational speech recognition," *arXiv preprint arXiv:1610.05256*, 2017.
- [4] M. Picheny, Z. Tüske, B. Kingsbury, K. Audhkhasi, X. Cui, and G. Saon, "Challenging the boundaries of speech recognition: The malach corpus," in *INTERSPEECH*, 2019.
- [5] F. Wessel and H. Ney, "Unsupervised training of acoustic models for large vocabulary continuous speech recognition," *IEEE Transactions on Speech and Audio Processing*, vol. 13, 2005.
- [6] T. Hori, R. Astudillo, T. Hayashi, Y. Zhang, S. Watanabe, and J. Le Roux, "Cycle-consistency training for end-to-end speech recognition," in *ICASSP*, 2019.

- [7] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *CoRR*, vol. abs/1807.03748, 2018.
- [8] A. Baevski, M. Auli, and A. Mohamed, "Effectiveness of self-supervised pre-training for speech recognition," *arXiv:1911.03912*, 2019.
- [9] L. Lamel, J.-L. Gauvain, and G. Adda, "Lightly supervised and unsupervised acoustic model training," *Computer Speech & Language*, 2002.
- [10] D. F. Harwath, A. Torralba, and J. R. Glass, "Unsupervised learning of spoken language with visual context," in *NeurIPS*, 2016.
- [11] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, "Visually grounded learning of keyword prediction from untranscribed speech," in *Interspeech*, 2017.
- [12] K. Singh, D. Okhonko, J. Liu, Y. Wang, F. Zhang, R. Girshick, S. Edunov, F. Peng, Y. Saraf, G. Zweig, and A. Mohamed, "Training asr models by generation of contextual information," *arXiv preprint arXiv:1910.12367*, 2019.
- [13] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *INTERSPEECH*, 2015.
- [14] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.
- [15] S. Krishnan Parthasarathi and N. Strom, "Lessons from building acoustic models with a million hours of speech," in *ICASSP*, 2019.
- [16] B. Li, R. Pang, T. Sainath, and Z. Wu, "Semi-supervised training for end-to-end models via weak distillation," in *ICASSP*, 2019.
- [17] T. Kudo and J. Richardson, "SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in *EMNLP*, 2018.
- [18] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*, 2016.
- [19] J. C. et. al., "Attention-based models for speech recognition," in *NeurIPS*, 2015.
- [20] T. Kemp and A. Waibel, "Unsupervised training of a speech recognizer: Recent experiments," *Eurospeech*, 1999.
- [21] C. Buciluundefined, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD*, 2006.
- [22] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [23] J. H. Wong and M. Gales, "Sequence student-teacher training of deep neural networks," in *Interspeech*, 2016.
- [24] H. K. Ryoichi Takashima, Li Sheng, "Investigation of sequence-level knowledge distillation methods for ctc acoustic models," in *ICASSP*, 2019.
- [25] N. I. Mun'im, Raden Mu'az and K. Shinoda, "Sequence-level knowledge distillation for model compression of attention-based sequence-to-sequence speech recognition," in *ICASSP*, 2019.
- [26] R. Masumura, M. Ithori, A. Takashima, T. Moriya, A. Ando, and Y. Shinohara, "Sequence-level consistency training for semi-supervised end-to-end automatic speech recognition," in *ICASSP*, 2020.
- [27] J. Ma, S. Matsoukas, O. Kimball, and R. Schwartz, "Unsupervised training on large amounts of broadcast news data," in *ICASSP*, 2006.
- [28] J. Gu, J. Bradbury, C. Xiong, V. O. Li, and R. Socher, "Non-autoregressive neural machine translation," *arXiv preprint arXiv:1711.02281*, 2017.
- [29] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," *ICASSP*, 2018.
- [30] A. Mohamed, D. Okhonko, and L. Zettlemoyer, "Transformers with convolutional context for ASR," *CoRR*, vol. abs/1904.11660, 2019.
- [31] S. K. et. al., "A comparative study on transformer vs RNN in speech applications," *CoRR*, vol. abs/1909.06317, 2019.
- [32] Y. Wang, A. Mohamed, D. Le, C. Liu, A. Xiao et al., "Transformer-based acoustic modeling for hybrid speech recognition," in *ICASSP*, 2020.
- [33] D. Le, X. Zhang, W. Zheng, C. Fügen, G. Zweig, and M. L. Seltzer, "From senones to chenones: Tied context-dependent graphemes for hybrid speech recognition," *ASRU*, 2019.
- [34] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2014.
- [35] 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.
- [36] K. Chen and Q. Huo, "Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering," in *ICASSP-2016*, March 2016.
- [37] D. Povey and B. Kingsbury, "Evaluation of proposed modifications to MPE for large scale discriminative training," in *ICASSP*, 2007.
- [38] V. Manohar, H. Hadian, D. Povey, and S. Khudanpur, "Semi-supervised training of acoustic models using Lattice-Free MMI," in *ICASSP*, 2018.
- [39] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," *arXiv:1911.04252*, 2019.
- [40] G. Zavaliagos and T. Colthurst, "Utilizing untranscribed training data to improve performance," in *DARPA Broadcast News Transcription and Understanding Workshop*, 1998.
- [41] A. J. et. al., "A summary of the 2012 JHU CLSP workshop on zero resource speech technologies and models of early language acquisition," in *ICASSP*, 2013.
- [42] H. Liao, E. McDermott, and A. Senior, "Large scale deep neural network acoustic modeling with semi-supervised training data for youtube video transcription," in *ASRU*, 2013.
- [43] B. Ramabhadran, "Exploiting large quantities of spontaneous speech for unsupervised training of acoustic models," in *Interspeech*, 2005, pp. 1617-1620.
- [44] P. Placeway and J. Lafferty, "Cheating with imperfect transcripts," in *ICSLP*, 1996.
- [45] H. Y. Chan and P. Woodland, "Improving broadcast news transcription by lightly supervised discriminative training," in *ICASSP*, 2004.
- [46] L. Wang, M. J. F. Gales, and P. C. Woodland, "Unsupervised training for mandarin broadcast news and conversation transcription," in *ICASSP*, 2007.
- [47] L. J. Ba and R. Caruana, "Do deep nets really need to be deep?" in *NIPS*. Cambridge, MA, USA: MIT Press, 2014.
- [48] V. Manohar, D. Povey, and S. Khudanpur, "Semi-supervised maximum mutual information training of deep neural network acoustic models," in *INTERSPEECH*, 2015.
- [49] J. Kahn, A. Lee, and A. Hannun, "Self-training for end-to-end speech recognition," *arXiv preprint arXiv:1909.09116*, 2019.
- [50] W.-N. Hsu, A. Lee, G. Synnaeve, and A. Hannun, "Semi-supervised speech recognition via local prior matching," 2020.
- [51] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *ICASSP*, 2015.
- [52] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko et al., "End-to-end asr: from supervised to semi-supervised learning with modern architectures," 2019.
- [53] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu et al., "Libri-light: A benchmark for asr with limited or no supervision," in *ICASSP*, 2020.