LETTER-BASED SPEECH RECOGNITION WITH GATED CONVNETS

Vitaliy Liptchinsky
Facebook AI Research
vitaliy888@fb.com

Gabriel Synnaeve
Facebook AI Research
gab@fb.com

Ronan Collobert
Facebook AI Research
locronan@fb.com

February 19, 2019

ABSTRACT

In the recent literature, "end-to-end" speech systems often refer to letter-based acoustic models trained in a sequence-to-sequence manner, either via a recurrent model or via a structured output learning approach (such as CTC (Graves et al., 2006)). In contrast to traditional phone (or senone)-based approaches, these "end-to-end" approaches alleviate the need of word pronunciation modeling, and do not require a "forced alignment" step at training time. Phone-based approaches remain however state of the art on classical benchmarks. In this paper, we propose a letter-based speech recognition system, leveraging a ConvNet acoustic model. Key ingredients of the ConvNet are Gated Linear Units and high dropout. The ConvNet is trained to map audio sequences to their corresponding letter transcriptions, either via a classical CTC approach, or via a recent variant called ASG (Collobert et al., 2016). Coupled with a simple decoder at inference time, our system matches the best existing letter-based systems on WSJ (in word error rate), and shows near state of the art performance on LibriSpeech (Panayotov et al., 2015).

1 Introduction

State of the art speech recognition systems leverage pronunciation models as well as speaker adaptation techniques involving speaker-specific features. These systems rely on lexicon dictionaries, which decompose words into one or more sequences of phones. Phones themselves are decomposed into smaller sub-word units, called senones. Senones are carefully selected through a procedure involving a phonetic-context-based decision tree built from another GMM/HMM system. In the recent literature, "end-to-end" speech systems attempt to break away from these hardcoded a-priori, the underlying assumption being that with enough data pronunciations should be implicitly inferred by the model, and speaker robustness should be also achieved. A number of works have thus naturally proposed ways how to learn to map audio sequences directly to their corresponding letter sequences. Recurrent models, structured-output learning or combination of both are the main contenders.

In this paper, we show that simple convolutional neural networks (CNNs) coupled with structured-output learning can outperform existing letter-based solutions. Our CNNs employ Gated Linear Units (GLU). Gated ConvNets have been shown to reduce the vanishing gradient problem, as they provide a linear path for the gradients while retaining non-linear capabilities, leading to state of the art performance both in natural language modeling and machine translation tasks (Dauphin et al., 2017; Gehring et al., 2017). We train our system with a structured-output learning approach, either with CTC (Graves et al., 2006) or ASG (Collobert et al., 2016). Coupled with a custom-made simple beam-search decoder, we exhibit word error rate (WER) performance matching the best existing letter-based systems, both for the WSJ and LibriSpeech datasets (Panayotov et al., 2015). While phone-based systems still lead on WSJ (81h of labeled data), our system is competitive with the existing state of the art systems on LibriSpeech (960h).

The rest of the paper is structured as follows: the next section goes over the history of the work in the automatic speech recognition area. We then detail the convolutional networks used for acoustic modeling, along with the structured-output learning and decoding approaches. The last section shows experimental results on WSJ and LibriSpeech.
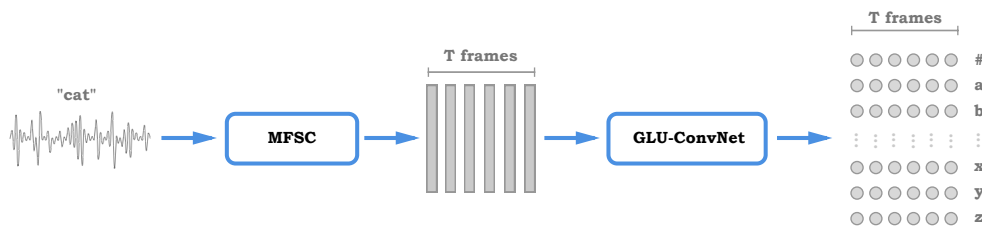
Figure 1: Overview of our acoustic model, which computes log-mel filterbanks (MFSC) that are fed to a Gated ConvNet. The ConvNet outputs one score for each letter in the dictionary, and for each input feature frame. At inference time, these scores are fed to a decoder (see Section 3.3) to form the most likely sequence of words. At training time, the scores are fed to the CTC or ASG criterions (see Figure 2) which promote sequences of letters leading to the target transcription sequence (here "c a t").

2 Background

The historic pipeline for speech recognition requires first training an HMM/GMM model to force align the units on which the final acoustic model operates (most often context-dependent phone or senone states) Woodland and Young (1993). The performance improvements brought by deep neural networks (DNNs) (Mohamed et al., 2012; Hinton et al., 2012) and convolutional neural networks (CNNs) (Sercu et al., 2016; Soltau et al., 2014) for acoustic modeling only extend this training pipeline. Current state of the art models on LibriSpeech also employ this approach (Panayotov et al., 2015; Peddinti et al., 2015b), with an additional step of speaker adaptation (Saon et al., 2013; Peddinti et al., 2015a). Departing from this historic pipeline, Senior et al. (2014) proposed GMM-free training, but the approach still requires to generate a forced alignment. Recently, maximum mutual information (MMI) estimation (Bahl et al., 1986) was used to train neural network acoustic models (Povey et al., 2016). The MMI criterion (Bahl et al., 1986) maximizes the mutual information between the acoustic sequence and word sequences or the Minimum Bayes Risk (MBR) criterion (Gibson and Hain, 2006), and belongs to segmental discriminative training criterions, although compatible with generative models.

Even though connectionist approaches (Lee et al., 1995; LeCun et al., 1995) long coexisted with HMM-based approaches, they had a recent resurgence. A modern work that directly cut ties with the HMM/GMM pipeline used a recurrent neural network (RNN) (Graves et al., 2013) for phoneme transcription with the connectionist temporal classification (CTC) sequence loss (Graves et al., 2006). This approach was then extended to character-based systems (Graves and Jaitly, 2014) and improved with attention mechanisms (Bahdanau et al., 2016; Chan et al., 2016). But the best such systems are often still behind state of the art phone-based (or senone-based) systems. Competitive end-to-end approaches leverage acoustic models (often ConvNet-based) topped with RNN layers as in (Hannun et al., 2014a; Miao et al., 2015; Saon et al., 2015; Amodei et al., 2016; Zhou et al., 2018; Zeyer et al., 2018) (e.g. a state of the art on WSJ (Chan and Lane, 2015a)), trained with a sequence criterion (the most popular ones being CTC (Graves et al., 2006) and MMI (Bahl et al., 1986)). A survey of segmental models can be found in (Tang et al., 2017). On conversational speech (that is not the topic of this paper), the state of the art is still held by complex ConvNets+RNNs acoustic models (which are also trained or refined with a sequence criterion), coupled with domain-adapted language models (Xiong et al., 2017; Saon et al., 2017).

3 Architecture

Our acoustic model (see an overview in Figure 1) is a Convolutional Neural Network (ConvNet) (LeCun et al., 1995), with Gated Linear Units (GLUs) (Dauphin et al., 2017) and dropout applied to activations of each layer except the output one. The model is fed with log-mel filterbank features, and is trained with either the Connectionist Temporal Classification (CTC) criterion (Graves et al., 2006), or with the ASG criterion: a variant of CTC that does not have blank labels but employs a simple duration model through letter transition scores (Collobert et al., 2016). At inference, the acoustic model is coupled with a decoder which performs a beam search, constrained with a count-based language model. We detail each of these components in the following.

3.1 Gated ConvNets for Acoustic Modeling

The acoustic model architecture is a 1D Gated Convolutional Neural Network (Gated ConvNet), trained to map a sequence of audio features to its corresponding letter transcription. Given a dictionary of letters $\mathcal{L}$, the ConvNet (which acts as a sliding-approach over the input sequence) outputs one score for each letter in the dictionary, for each input frame. In the transcription, words are separated by a special letter, denoted #.

1D ConvNets were introduced early in the speech community, and are also referred as Time-Delay Neural Networks (TDNNs) (Waibel, 1989). Gated ConvNets (Dauphin et al., 2017) stack 1D convolutions with Gated Linear Units. More formally, given an input sequence $\mathbf{X} \in \mathbb{R}^{T \times d^i}$ with T frames of d-dimensional vectors, the i^{th} layer of our network performs the following computation:

$$h^i(\mathbf{X}) = (\mathbf{X} * \mathbf{W}^i + \mathbf{b}^i) \otimes \sigma(\mathbf{X} * \mathbf{V}^i + \mathbf{c}^i), \qquad (1)$$

where $*$ is the convolution operator, $\mathbf{W}^i$, $\mathbf{V}^i \in \mathbb{R}^{d^{i+1} \times d^i \times k^i}$ and $\mathbf{b}^i$, $\mathbf{c}^i \in \mathbb{R}^{d^{i+1}}$ are the learned parameters (with convolution kernel size k^i), $\sigma(\cdot)$ is the sigmoid function and $\otimes$ is the element-wise product between matrices.

Gated ConvNets have been shown to reduce the vanishing gradient problem, as they provide a linear path for the gradients while retaining non-linear capabilities, leading to state of the art performance both for natural language modeling and machine translation tasks (Dauphin et al., 2017; Gehring et al., 2017).

3.1.1 Feature Normalization and Zero-Padding

Each input feature sequence is normalized to have zero mean and unit variance. Given an input sequence $\mathbf{X} \in \mathbb{R}^{T \times d}$, a convolution with kernel size k will output $T - k + 1$ frames, due to border effects. To compensate those border effects, we pad the log-mel filterbanks $\mathbf{X}^0$ with zeroed frames. To take into account the whole network, the padding size is $\sum_i (k^i - 1)$, divided into two equal parts at the beginning and the end of the sequence.

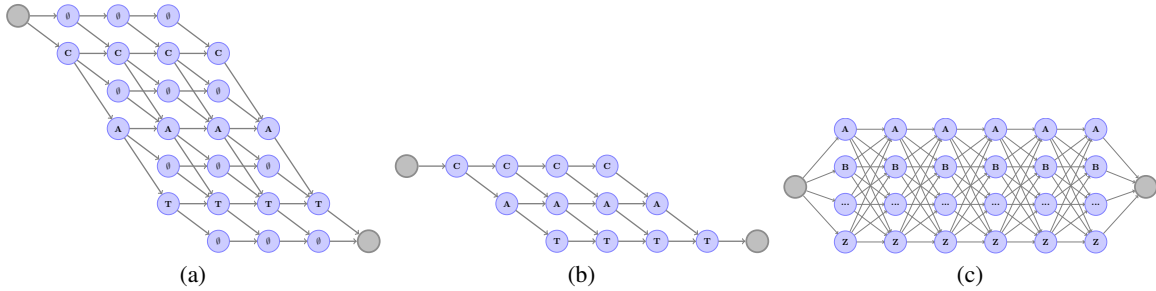### 3.2 Acoustic Model Training

Figure 2: (a) The CTC graph which represents all the acceptable sequences of letters for the transcription "cat" over 6 frames. (b) The same graph used by ASG, where blank labels have been discarded. (c) The fully connected graph describing all possible sequences of letter; this graph is used for normalization purposes in ASG. Un-normalized transitions scores are possible on edges of these graphs. At each time step, nodes are assigned a conditional un-normalized score, output by the Gated ConvNet acoustic model.

We considered two structured-output learning approaches to train our acoustic models: the Connectionist Temporal Classification (CTC), and a variant called AutoSeG (ASG).

3.2.1 The CTC Criterion

CTC (Graves et al., 2006) efficiently enumerates all possible sequences of sub-word units (e.g. letters) which can lead to the correct transcription, and promotes the score of these sequences. CTC also allows a special "blank" state to be optionally inserted between each sub-word unit. The rationale behind the blank state is two-fold: (i) modeling "garbage" frames which might occur between each letter and (ii) identifying the separation between two identical consecutive sub-word units in a transcription. Figure 2a shows the CTC graph describing all the possible sequences of letters leading to the word "cat", over 6 frames. We denote $\mathcal{G}_{ctc}(\theta, T)$ the CTC acceptance graph over T frames for a given transcription θ, and $\pi = \pi_1, \ldots, \pi_T \in \mathcal{G}_{ctc}(\theta, T)$ a path in this graph representing a (valid) sequence of letters for this transcription. CTC assumes that the network outputs probability scores, normalized at the frame level. At each time step t, each node of the graph is assigned with its corresponding log-probability letter i (that we denote $f_i^t(\mathbf{X})$) output by the acoustic model (given an acoustic sequence $\mathbf{X}$). CTC minimizes the Forward score over the graph $\mathcal{G}_{ctc}(\theta, T)$:

$$CTC(\theta, T) = - \operatorname*{logadd}_{\pi \in \mathcal{G}_{ctc}(\theta, T)} \sum_{t=1}^{T} f_{\pi_t}^t(\mathbf{X}), \qquad (2)$$

where the "logadd" operation (also called "log-sum-exp") is defined as $\operatorname{logadd}(a, b) = \log(\exp(a) + \exp(b))$. This overall score can be efficiently computed with the Forward algorithm.

3.2.2 The ASG Criterion

Blank labels introduce code complexity when decoding letters into words. Indeed, with blank labels "ø", a word gets many entries in the sub-word unit transcription dictionary (e.g. the word "cat" can be represented as "c a t", "c ø a t", "c ø a t", "c ø a ø t", etc... – instead of only "c a t"). We replace the blank label by special letters modeling repetitions of preceding letters. For example "caterpillar" can be written as "caterpil1ar", where "1" is a label to represent one repetition of the previous letter.

The AutoSeG (ASG) criterion (Collobert et al., 2016) removes the blank labels from the CTC acceptance graph $\mathcal{G}_{ctc}(\theta, T)$ (shown in Figure 2a) leading to a simpler graph that we denote $\mathcal{G}_{asg}(\theta, T)$ (shown in Figure 2b). In contrast to CTC which assumes per-frame normalization for the acoustic model scores, ASG implements a sequence-level normalization to prevent the model from diverging (the corresponding graph enumerating all possible sequences of letters is denoted $\mathcal{G}_{asg}(\theta, T)$, as shown in Figure 2c). ASG also uses unnormalized transition scores $g_{i,j}(\cdot)$ on each edge of the graph, when moving from label i to label j, that are trained jointly with the acoustic model. This leads to the following criterion::

$$
\begin{aligned}
ASG(\theta, T) = - &\operatorname*{logadd}_{\pi \in \mathcal{G}_{asg}(\theta, T)} \sum_{t=1}^{T} (f_{\pi_t}^t(\mathbf{X}) + g_{\pi_{t-1}, \pi_t}(\mathbf{X})) \\
+ &\operatorname*{logadd}_{\pi \in \mathcal{G}_{full}(\theta, T)} \sum_{t=1}^{T} (f_{\pi_t}^t(\mathbf{X}) + g_{\pi_{t-1}, \pi_t}(\mathbf{X})) \,.
\end{aligned}
\tag{3}
$$

The left-hand part in Equation (3) promotes the score of letter sequences leading to the right transcription (as in Equation (2) for CTC), and the right-hand part demotes the score of all sequences of letters. As for CTC, these two parts can be efficiently computed with the Forward algorithm.

When removing transitions in Equation (3), the sequence-level normalization becomes equivalent to the frame-level normalization found in CTC, and the ASG criterion is mathematically equivalent to CTC with no blank labels. However, in practice, we observed that acoustic models trained with a transition-free ASG criterion had a hard time to converge.

3.2.3 Other Training Considerations

We apply dropout at the output to all layers of the acoustic model. Dropout retains each output with a probability p, by applying a multiplication with a Bernoulli random variable taking value 1 with probability p and 0 otherwise (Srivastava et al., 2014).

Following the original implementation of Gated ConvNets (Dauphin et al., 2017), we found that using both weight normalization (Salimans and Kingma, 2016) and gradient clipping (Pascanu et al., 2013) were speeding up training convergence. The clipping we implemented performs:

$$
\widetilde{\nabla} C = \max(||\nabla C||, \epsilon) \frac{\nabla C}{||\nabla C||} \,,
\tag{4}
$$

where C is either the CTC or ASG criterion, and ϵ is some hyper-parameter which controls the maximum amplitude of the gradients.

3.3 Beam-Search Decoder

We wrote our own one-pass decoder, which performs a simple beam-search with beam thresholding, histogram pruning and language model smearing (Steinbiss et al., 1994). We kept the decoder as simple as possible (under 1000 lines of C code). We did not implement any sort of model adaptation before decoding, nor any word graph rescoring. Our decoder relies on KenLM (Heafield et al., 2013) for the language modeling part. It also accepts unnormalized acoustic scores (transitions and emissions from the acoustic model) as input. The decoder attempts to maximize the following:

$$
\begin{aligned}
\mathcal{L}(\theta) = &\operatorname*{logadd}_{\pi \in \mathcal{G}_{lex}(\theta, T)} \sum_{t=1}^{T} (f_{\pi_t}(x) + g_{\pi_{t-1}, \pi_t}(x)) \\
&+ \alpha \log P_{lm}(\theta) + \gamma |\{i | \pi_i = \texttt{\#}\}| \,,
\end{aligned}
\tag{5}
$$

where $\mathcal{G}_{lex}(\theta, T)$ is a graph constrained by lexicon, $P_{lm}(\theta)$ is the probability of the language model given a transcription θ, α, β, and γ are three hyper-parameters which control the weight of the language model, and the silence (#) insertion penalty, respectively.

4

Table 1: Architecture details. "#conv." is the number of convolutional layers. Dropout amplitude, "#hu" (number of output hidden units) and "kw" (convolution kernel width) are provided for the first and last layer (all are linearly increased with layer depth). The size of the final layer is also provided.

Dataset	Architecture	#conv.	dropout first/last layer	#hu first/last layer	kw first/last layer	#hu full connect
WSJ	Low Dropout	17	0.25/0.25	100/375	3/21	1000
LibriSpeech	Low Dropout	17	0.25/0.25	200/750	13/27	1500
	High Dropout	19	0.20/0.60	200/1000	13/29	2000

The beam of the decoder tracks paths with highest scores according to Equation (5), by bookkeeping pairs of (language model, lexicon) states, as it goes through time. The language model state corresponds to the $(n-1)$-gram history of the n-gram language model, while the lexicon state is the sub-word unit position in the current word hypothesis. To maintain diversity in the beam, paths with identical (language model, lexicon) states are merged. Note that traditional decoders combine the scores of the merged paths with a $\max(\cdot)$ operation (as in a Viterbi beam-search) – which would correspond to a $\max(\cdot)$ operation in Equation (5) instead of $\mathrm{logadd}(\cdot)$. We consider instead the $\mathrm{logadd}(\cdot)$ operation (as first suggested by Bottou (1991)), as it takes into account the contribution of all the paths leading to the same transcription, in the same way we do during training (see Equation (3)). In Section 4.1, we show that this leads to better accuracy in practice.

4 Experiments

We benchmarked our system on WSJ (about 81h of labeled audio data) and LibriSpeech (Panayotov et al., 2015) (about 960h). We kept the original 16 kHz sampling rate. For WSJ, we considered the classical setup SI284 for training, DEV93 for validation, and EVAL92 for evaluation. For LibriSpeech, we considered the two available setups CLEAN and OTHER. All the hyper-parameters of our system were tuned on validation sets. Test sets were used only for the final evaluations.

The letter vocabulary $\mathcal{L}$ contains 30 graphemes: the standard English alphabet plus the apostrophe, silence (#), and two special "repetition" graphemes which encode the duplication (once or twice) of the previous letter (see Section 3.2.2). Decoding is achieved with our own decoder (see Section 3.3). We used standard language models for both datasets, *i.e.* a 4-gram model (with about $165K$ words in the dictionary) trained on the provided data for WSJ, and a 4-gram model[1] (about $200K$ words) for LibriSpeech. In the following, we either report letter-error-rates (LERs) or word-error-rates (WERs).

Training was performed with stochastic gradient descent on WSJ, and mini-batches of 4 utterances on LibriSpeech. Clipping parameter (see Equation (4)) was set to $\epsilon = 0.2$. We used a momentum of 0.9. Input features, log-mel filterbanks, were computed with 40 coefficients, a 25 ms sliding window and 10 ms stride.

We implemented everything using TORCH7[2]. The CTC and ASG criterions, as well as the decoder were implemented in C (and then interfaced into TORCH).

4.1 Architecture

We tuned our acoustic model architectures by grid search, validating on the dev sets. We consider here two architectures, with low and high amount of dropout (see the parameter p in Section 3.2.3). Table 1 reports the details of our architectures. The amount of dropout, number of hidden units, as well as the convolution kernel width are increased linearly with the depth of the neural network. Note that as we use Gated Linear Units (see Section 3.1), each layer is duplicated as stated in Equation (1). Convolutions are followed by a fully connected layer, before the final layer which outputs 30 scores (one for each letter in the dictionary). Concerning WSJ, the LOW DROPOUT ($p = 0.2$) architecture has about 17M trainable parameters. For LibriSpeech, architectures have about 130M and 208M of parameters for the LOW DROPOUT ($p = 0.2$) and HIGH DROPOUT ($p = 0.2 \rightarrow 0.6$) architectures, respectively.

---

[1] http://www.openslr.org/11.
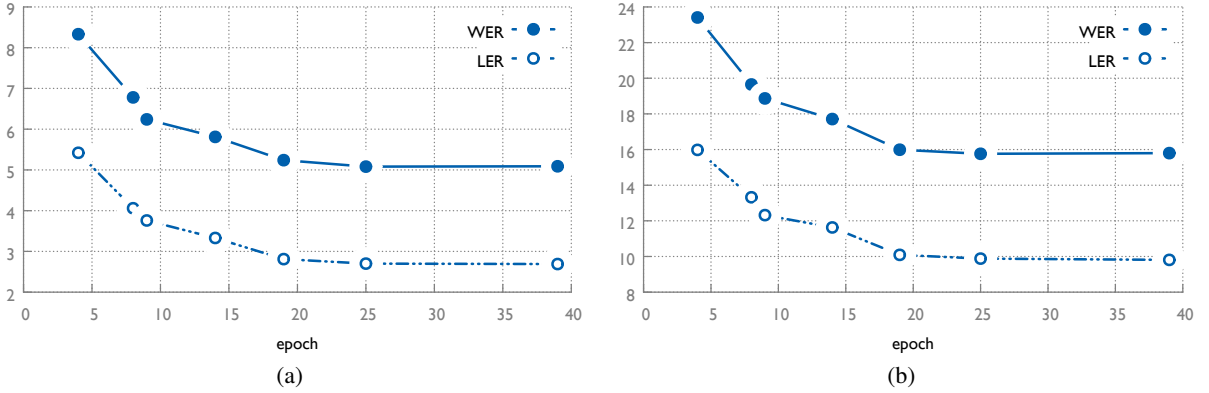[2] http://www.torch.ch.

Figure 3: LibriSpeech Letter Error Rate (LER) and Word Error Rate (WER) for the first training epochs of our LOW DROPOUT architecture. (a) is on `dev-clean`, (b) on `dev-other`.

Table 2: Comparison in LER and WER of variants of our model on (a) WSJ and (b) LibriSpeech. LER is computed with *no* decoding. Operator max and logadd refer to the aggregation of beam hypotheses (see Section 3.3).

(a) WSJ

model	dev93 LER	dev93 WER
LOW DROP. (ASG, max)	–	10.0
LOW DROP. (ASG, logadd)	7.2	9.8

(b) LibriSpeech

model	dev-clean LER	dev-clean WER	dev-other LER	dev-other WER
LOW DROP. (ASG, logadd)	2.7	4.8	9.8	15.2
HIGH DROP. (CTC)	2.6	4.7	9.5	14.9
HIGH DROP. (ASG, max)	–	4.7	–	14.0
HIGH DROP. (ASG, logadd)	2.3	4.6	9.0	13.8

4.1.1 Analysis

Figure 3 shows the LER and WER on the LibriSpeech development sets, for the first 40 training epochs of our LOW DROPOUT architecture. LER and WER appear surprisingly well correlated, both on the "clean" and "other" version of the dataset.

In Table 2b, we report WERs on the LibriSpeech development sets, both for our LOW DROPOUT and HIGH DROPOUT architectures. Increasing dropout regularize the acoustic model in a way which impacts significantly generalization, the effect being stronger on noisy speech.

Table 2a and Table 2b also report the WER for the decoder ran with the $\max(\cdot)$ operation (instead of $\mathrm{logadd}(\cdot)$ for other results) used to aggregate paths in the beam with identical (language model, lexicon) states. It appears advantageous (as there is no code complexity increase in the decoder – one only needs to replace $\max(\cdot)$ by $\mathrm{logadd}(\cdot)$ in the code) to use the $\mathrm{logadd}(\cdot)$ aggregation.
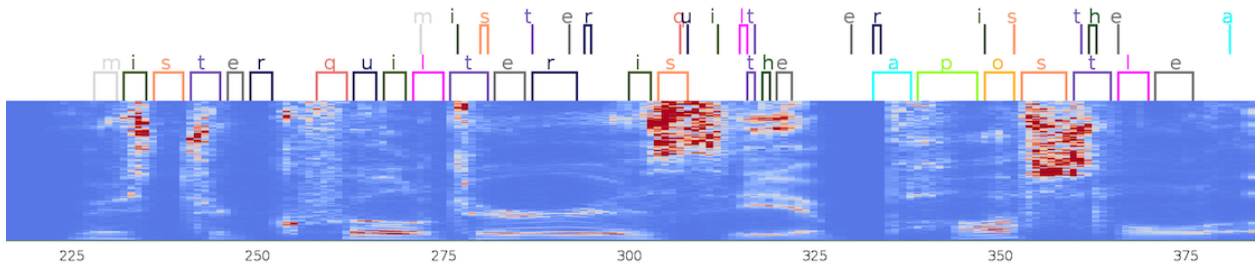

Figure 4: Comparison of alignments produced by the models with CTC (top) and ASG (bottom) criterions on audio spectrogram over time (each time frame on *X* axis corresponds to a 40ms window with 10 ms stride).

Table 3: Comparison of different near state of the art ASR systems on LibriSpeech. We report the type of acoustic model used for various systems, as well as the type of sub-word units. HMM stands for Hidden Markov Model, CNN for ConvNet; when not specified, CNNs are 1D. pNorm is a particular non-linearity (Zhang et al., 2014). We also report extra information (besides word transcriptions) which might be used by each system, including speaker adaptation, or any other domain-specific data.

Paper	Acoustic Model	Sub-word	Spkr Adapt.	Extra Resources
Panayotov et al. (2015)	HMM+DNN+pNorm	phone	fMLLR	phone lexicon
Peddinti et al. (2015b)	HMM+CNN	phone	iVectors	phone lexicon
Povey et al. (2016)	HMM+CNN	phone	iVectors	phone lexicon, phone LM, data augm.
Ko et al. (2015)	HMM+CNN+pNorm	phone	iVectors	phone lexicon, data augm.
Amodei et al. (2016)	2D-CNN+RNN	letter	*none*	11.9Kh train set, Common Crawl LM
Zhou et al. (2018)	CNN+GRU+policy learning	letter	*none*	data augmentation
Zeyer et al. (2018)	RNN+attention	letter	*none*	LSTM LM
this paper	GLU-CNN	letter	*none*	*none*

Figure 4 depicts alignments of the models with CTC and ASG criterions when forced aligned to a given target. Our analysis shows that the model with CTC criterion exhibits 500 ms delay compared to the model with ASG criterion. Similar observation was also previously noted in Sak et al. (2015).

4.2 Comparison with other systems

In Table 4, we compare our system with existing phone-based and letter-based approaches on WSJ and LibriSpeech. Phone-based acoustic state of the art models are reported as reference. These systems output in general senones; senones are carefully selected through a procedure involving a phonetic-context-based decision tree built from another GMM/HMM system. Phone-based systems also require an additional word lexicon which translates words into a sequence of phones. Most state of the art systems also perform speaker adaptation; iVectors compute a speaker embedding capturing both speaker and environment information (Xue et al., 2014), while fMMLR is a two-pass decoder technique which computes a speaker transform in the first pass (Gales and Woodland, 1996). Even though Table 3 associates speaker adaptation exclusively with phone-based systems, speaker adaptation can be also applied to letter-based systems.

State of the art performance for letter-based models on LibriSpeech is held by DEEP SPEECH 2 (Amodei et al., 2016) and (Zeyer et al., 2018) on noisy and clean subsets respectively. On WSJ state of the art performance is held by DEEP SPEECH 2. DEEP SPEECH 2 uses an acoustic model composed of a ConvNet and a Recurrent Neural Network (RNN). DEEP SPEECH 2 relies on a lot of extra speech data at training, combined with a very large 5-gram language model at inference time to make the letter-based approach competitive. Our system outperforms DEEP SPEECH 2 on clean data, even though our system has been trained with an order of magnitude less data. Acoustic model in (Zeyer et al., 2018) is also based on RNNs and in addition employs attention mechanism. With LSTM language model their system shows lower WER than our, but with a simple 4-gram language model our system has slightly lower WER.

On WSJ the state of the art is a phone-based approach (Chan and Lane, 2015b) which leverages an acoustic model combining CNNs, bidirectional LSTMs, and deep fully connected neural networks. The system also performs speaker adaptation at inference. We also compare with existing letter-based approaches on WSJ, which are abundant in the literature. They rely on recurrent neural networks, often bi-directional, and in certain cases combined with ConvNet architectures. Our system matches the best reported letter-based WSJ performance. The Gated ConvNet appears to be very strong at modeling complete words as it achieves 6.7% WER on LibriSpeech clean data even with no decoder, i.e. on the raw output of the neural network.

Concerning LibriSpeech, we summarize existing state of the art systems in Table 3. We highlighted the acoustic model architectures, as well as the type of underlying sub-word units.

5 Conclusion

We have introduced a simple end-to-end automatic speech recognition system, which combines a ConvNet acoustic model with Gated Linear Units, and a simple beam-search decoder. The acoustic model is trained to map audio

Table 4: Comparison in WER of our model with other systems on WSJ and LibriSpeech. Systems with ⋆ or † use additional data or data augmentation at training, respectively.

	WSJ eval92	LibriSpeech test-clean	LibriSpeech test-other
(Povey et al., 2016)	4.3	-	-
(Panayotov et al., 2015)	3.9	5.5	14.0
(Peddinti et al., 2015b)	-	4.8	-
(Chan and Lane, 2015b)	3.5	-	-
(Ko et al., 2015)†	-	-	12.5
(Hannun et al., 2014b)	14.1	-	-
(Bahdanau et al., 2016)	9.3	-	-
(Graves and Jaitly, 2014)	8.2	-	-
(Miao et al., 2015)	7.3	-	-
(Chorowski and Jaitly, 2016)	6.7	-	-
(Liu et al., 2017)†	6.7	-	-
(Hori et al., 2017)	5.6	-	-
(Zhou et al., 2018)†	5.5	5.7	15.2
(Amodei et al., 2016)⋆	3.6	5.3	13.3
(Zeyer et al., 2018)	-	4.8	14.7
(Zeyer et al., 2018) (*LSTM LM*)	-	3.8	12.8
this paper (CTC)	-	5.1	16.0
this paper (ASG)	5.6	4.8	14.5

sequences to sequences of characters using a structured-output learning approach based on a variant of CTC. Our system outperforms existing letter-based approaches (which do not use extra data at training time or powerful LSTM language models), both on WSJ and LibriSpeech. Overall phone-based approaches are still holding the state of the art, but our system's performance is competitive on LibriSpeech, suggesting pronunciations is implicitly well modeled with enough training data. Further work should include leveraging speaker identity, training from the raw waveform, data augmentation, training with more data, and better language models.

References

D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen, J. Chen, J. Chen, Z. Chen, M. Chrzanowski, A. Coates, G. Diamos, et al. Deep speech 2 : End-to-end speech recognition in english and mandarin. In *International Conference on Machine Learning (ICML)*, pages 173–182, 2016.

D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio. End-to-end attention-based large vocabulary speech recognition. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4945–4949, 2016.

L. R. Bahl, P. F. Brown, P. V. de Souza, and R. L. Mercer. Maximum mutual information estimation of hidden Markov model parameters for speech recognition. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 49–52, 1986.

L. Bottou. *Une Approche théorique de l'Apprentissage Connexionniste: Applications à la Reconnaissance de la Parole*. PhD thesis, 1991.

W. Chan and I. Lane. Deep recurrent neural networks for acoustic modelling. *arXiv preprint arXiv:1504.01482*, 2015a.

W. Chan and I. Lane. Deep recurrent neural networks for acoustic modelling. *arXiv:1504.01482*, 2015b.

W. Chan, N. Jaitly, Q. Le, and O. Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964, 2016.

J. Chorowski and N. Jaitly. Towards better decoding and language model integration in sequence to sequence models. *arXiv:1612.02695*, 2016.

R. Collobert, C. Puhrsch, and G. Synnaeve. Wav2letter: an end-to-end convnet-based speech recognition system. *arXiv:1609.03193*, 2016.

Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier. Language modeling with gated convolutional nets. In *International Conference on Machine Learning (ICML)*, 2017.

M. J. F. Gales and P. C. Woodland. Mean and variance adaptation within the MLLR framework. *Computer Speech and Language*, 10(4):249–264, 1996.

J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin. Convolutional sequence to sequence learning. In *International Conference on Machine Learning (ICML)*, 2017.

M. Gibson and T. Hain. Hypothesis spaces for minimum Bayes risk training in large vocabulary speech recognition. In *Interspeech*, pages 2406—-2409, 2006.

A. Graves and N. Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *International Conference on Machine Learning*, pages 1764–1772, 2014.

A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *International Conference on Machine Learning (ICML)*, pages 369–376. ACM, 2006.

A. Graves, A.-R. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6645–6649, 2013.

A. Y. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *arXiv:1412.5567*, 2014a.

A. Y. Hannun, A. L. Maas, D. Jurafsky, and A. Y. Ng. First-pass large vocabulary continuous speech recognition using bi-directional recurrent dnns. *arXiv:1408.2873*, 2014b.

K. Heafield, I. Pouzyrevsky, J. H. Clark, and P. Koehn. Scalable modified kneser-ney language model estimation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 690–696, 2013.

G. Hinton, L. Deng, D. Yu, G. Dahl, A. rahman Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury. Deep neural networks for acoustic modeling in speech recognition. *Signal Processing Magazine*, 29(6):82–97, 2012.

T. Hori, S. Watanabe, and J. R. Hershey. Multi-level language modeling and decoding for open vocabulary end-to-end speech recognition. In *Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE*. IEEE, 2017.

T. Ko, V. Peddinti, D. Povey, and S. Khudanpur. Audio augmentation for speech recognition. In *Interspeech*, 2015.

Y. LeCun, Y. Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

T. Lee, P. Ching, and L.-W. Chan. An rnn based speech recognition system with discriminative training. In *Fourth European Conference on Speech Communication and Technology*, 1995.

H. Liu, Z. Zhu, X. Li, and S. Satheesh. Gram-ctc: Automatic unit selection and target decomposition for sequence labelling. *CoRR*, abs/1703.00096, 2017.

Y. Miao, M. Gowayyed, and F. Metze. Eesen: End-to-end speech recognition using deep RNN models and WFST-based decoding. In *Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2015.

A.-R. Mohamed, G. E. Dahl, and G. Hinton. Acoustic modeling using deep belief networks. *Transactions on Audio, Speech, and Language Processing*, 20(1):14–22, 2012.

V. Panayotov, G. Chen, D. Povey, and S. Khudanpur. Librispeech: an ASR corpus based on public domain audio books. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, 2015.

R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning (ICML)*, 2013.

V. Peddinti, G. Chen, V. Manohar, T. Ko, D. Povey, and S. Khudanpur. JHU ASpIRE system: Robust LVCSR with TDNNs, iVector adaptation, and RNN-LMs. In *Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2015a.

V. Peddinti, D. Povey, and S. Khudanpur. A time delay neural network architecture for efficient modeling of long temporal contexts. In *Interspeech*, 2015b.

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*, pages 2751–2755, 2016.

H. Sak, F. de Chaumont Quitry, T. Sainath, K. Rao, et al. Acoustic modelling with cd-ctc-smbr lstm rnns. In *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*, pages 604–609. IEEE, 2015.

T. Salimans and D. P. Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 901–909. 2016.

G. Saon, H. Soltau, D. Nahamoo, and M. Picheny. Speaker adaptation of neural network acoustic models using I-Vectors. In *Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 55–59, 2013.

G. Saon, H.-K. J. Kuo, S. Rennie, and M. Picheny. The IBM 2015 english conversational telephone speech recognition system. *arXiv:1505.05899*, 2015.

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:1703.02136*, 2017.

A. Senior, G. Heigold, M. Bacchiani, and H. Liao. GMM-free DNN training. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5639–5643, 2014.

T. Sercu, C. Puhrsch, B. Kingsbury, and Y. LeCun. Very deep multilingual convolutional neural networks for LVCSR. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4955–4959, 2016.

H. Soltau, G. Saon, and T. N. Sainath. Joint training of convolutional and non-convolutional neural networks. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5572–5576, 2014.

N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research (JMLR)*, 15(Jun):1929–1958, 2014.

V. Steinbiss, B.-H. Tran, and H. Ney. Improvements in beam search. In *International Conference on Spoken Language Processing (ICSLP)*, volume 94, pages 2143–2146, 1994.

H. Tang, L. Lu, L. Kong, K. Gimpel, K. Livescu, C. Dyer, N. A. Smith, and S. Renals. End-to-end neural segmental models for speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11:1254–1264, 2017.

A. Waibel. Modular construction of time-delay neural networks for speech recognition. *Neural Computation*, 1(1): 39–46, 1989.

P. C. Woodland and S. J. Young. The HTK tied-state continuous speech recogniser. In *Eurospeech*, 1993.

W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig. The Microsoft 2016 conversational speech recognition system. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5255–5259, 2017.

S. Xue, O. Abdel-Hamid, H. Jiang, L. Dai, and Q. Liu. Fast adaptation of deep neural network based on discriminant codes for speech recognition. *Transactions on Audio, Speech and Language Processing*, 22(12):1713–1725, 2014.

A. Zeyer, K. Irie, R. Schlüter, and H. Ney. Improved training of end-to-end attention models for speech recognition. *arXiv:1805.03294*, 2018.

X. Zhang, J. Trmal, D. Povey, and S. Khudanpur. Improving deep neural network acoustic models using generalized maxout networks. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 215–219, 2014.

Y. Zhou, C. Xiong, and R. Socher. Improving end-to-end speech recognition with policy learning. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.