
DONUT: CTC-based Query-by-Example Keyword Spotting

Loren Lugosch
Fluent.ai

loren.lugosch@fluent.ai

Samuel Myer
Fluent.ai

sam.myer@fluent.ai

Vikrant Singh Tomar
Fluent.ai

vikrant.tomar@fluent.ai

Abstract

Keyword spotting—or wakeword detection—is an essential feature for hands-free operation of modern voice-controlled devices. With such devices becoming ubiquitous, users might want to choose a personalized custom wakeword. In this work, we present DONUT, a CTC-based algorithm for online query-by-example keyword spotting that enables custom wakeword detection. The algorithm works by recording a small number of training examples from the user, generating a set of label sequence hypotheses from these training examples, and detecting the wakeword by aggregating the scores of all the hypotheses given a new audio recording. Our method combines the generalization and interpretability of CTC-based keyword spotting with the user-adaptation and convenience of a conventional query-by-example system. DONUT has low computational requirements and is well-suited for both learning and inference on embedded systems without requiring private user data to be uploaded to the cloud.

1 Introduction

Keyword spotting enables a user to activate a device conveniently, without pushing a button or otherwise physically touching the device, by detecting when the user speaks a certain “wakeword”. Most wakeword detectors use a pre-set wakeword (for example, “Hey Siri”, “OK Google”, or “Alexa”). It is desirable for users to be able to choose their own wakeword instead of using a pre-set wakeword. For instance, if the device is a pet robot, a custom wakeword would allow users to give their robot a name, which gives the device a more personal feel.

Neural networks can be trained to perform very accurate online keyword spotting for a pre-defined wakeword [1; 2]. However, to recognize a new wakeword, the network must be retrained, which requires time and many training examples. A more sample-efficient approach based on connectionist temporal classification (CTC) [3] overcomes this problem by representing the keyword as a sequence of phonetic labels and using the CTC forward algorithm to efficiently compute a score for the sequence given the neural network output [4; 5; 6]. Thus, to recognize a custom wakeword, the user can provide the label sequence corresponding to the desired phrase without retraining the network.

A disadvantage of conventional CTC-based keyword spotting is that it is “query-by-string”: that is, the desired wakeword must be provided in text form. Query-by-string keyword spotting is somewhat inconvenient for the user and requires that the device have a text interface in addition to the voice interface. Additionally, because the wakeword is provided through text, the wakeword model may not actually match the user’s own pronunciation of the phrase.

A more natural method for custom wakeword detection is “query-by-example” keyword spotting. In a query-by-example system, the user teaches the system the desired wakeword by recording a few training examples, and the keyword spotter uses some form of template matching to compare incoming audios with these training examples to detect the wakeword. In dynamic time warping (DTW)-based keyword spotting, for example, a variable-length sequence of feature vectors, such as Mel-filterbank cepstral coefficients (MFCCs) [7] or phoneme posteriors [8; 9; 10], is extracted from the query audio and test audio, and the DTW alignment score between query and test is used as the detection score. Other template-matching approaches compare fixed-length feature vectors, such as the final hidden states of a pre-trained recurrent neural network (RNN) [11] or the output of a Siamese network [12; 13], using the cosine distance.

Systems that use template matching are difficult to interpret, and therefore difficult to debug and optimize. For instance, it is hard to say why a keyword is incorrectly detected or not detected in a system based on dynamic time warping (DTW) simply by inspecting the DTW matrix. Likewise, the hidden states of RNNs can sometimes be interpreted (c.f. [14], [15]), but this is currently only possible with some luck and ingenuity. In contrast, a CTC-based model is easy to interpret. The wakeword model itself is interpretable: it consists simply of a human-readable string, like “ALEXA” or “AH L EH K S AH”, rather than a vector of real numbers. Inference is interpretable because the neural network outputs are peaky and sparse (the “blank” symbol has probability ≈ 1 at almost all timesteps), so it is easy to determine what the network “hears” for any given audio and whether it hears the labels of the wakeword [16]. This is a useful property because it enables the system designer to take corrective action. For instance, one might identify that a particular label is not well-recognized and augment the training data with examples of this label.

In this paper, we propose a new method for custom wakeword detection that combines the convenience and speaker-adaptive quality of query-by-example methods with the generalization power and interpretability of CTC-based keyword spotting. We call our method “DONUT”, since detection requires $O(NUT)$ operations given the neural network output, where N , U , and T are small numbers defined later in the paper. The method works as follows: the user records a few training examples of the keyword, and a beam search is used to estimate the labels of the keyword. The algorithm maintains an N -best list of label sequence hypotheses to minimize the error that may be incurred by incorrectly estimating the labels. At inference time, each hypothesis is scored using the forward algorithm, and the hypothesis scores are aggregated to obtain a single detection score.

In the rest of the paper, we describe the proposed method and show that it achieves good performance compared with other query-by-example methods, yet generates easily interpretable models and matches the user’s pronunciation better than when the label sequence is supplied through text.

2 Proposed method

This section describes the model, learning, and inference for DONUT (Fig. 1), as well as the memory, storage, and computational requirements.


2.1 Model

The proposed method uses a model composed of a wakeword model and a label model. Here we give more detail on these two components.

2.1.1 Wakeword model

We can model the user’s chosen wakeword as a sequence of labels $\mathbf{y} = \{y_u \in \mathcal{A} \mid u = 1, \dots, U\}$, where \mathcal{A} is the set of possible labels, and U is the length of the sequence. The labels could be phonemes, graphemes, or other linguistic subunits; in this work, we use phonemes. It is generally not possible to perfectly estimate \mathbf{y} from only a few training examples. Therefore, we maintain multiple hypotheses as to what the true sequence might be, along with a confidence for each hypothesis, and make use of all of these hypotheses during inference. A trained wakeword model thus consists of a set of label sequences and confidences.

Learning



Inference




Figure 1: Illustration of the proposed method. Here, two training examples are recorded by the user, with a beam search of width $B = 3$ and $N = 2$ hypotheses kept per training example.

2.1.2 Label model

The label model ϕ is a neural network trained using CTC on a speech corpus where each audio has a transcript of labels from the label set \mathcal{A} . The network accepts an audio in the form of a sequence of acoustic feature vectors $\mathbf{x} = \{x_t \in \mathbb{R}^d \mid t = 1, \dots, T\}$, where d is the number of features per frame, and T is the number of frames. The network outputs a posteriorgram $\boldsymbol{\pi} = f_\phi(\mathbf{x}) = \{\pi_t \in \mathbb{R}^{1+|\mathcal{A}|} \mid t = 1, \dots, T\}$ representing the posterior probabilities of each of the labels and the CTC “blank” symbol at each timestep.

2.2 Learning

Algorithm 1 describes the learning phase. The user records three examples of the wakephrase, here denoted by $\mathbf{x}^{train,1}$, $\mathbf{x}^{train,2}$, and $\mathbf{x}^{train,3}$. Once the user has recorded the audios, the label posteriors $\boldsymbol{\pi}^{train,i}$ for each audio are computed using the label model ϕ . The CTCBeamSearch function then runs a beam search of width B over the label posteriors and returns a list of B probable label sequences and their corresponding log probabilities. More details on the beam search algorithm for CTC models can be found in [17]. The top N hypotheses $\hat{\mathbf{y}}_{1,\dots,N}^{train,i}$ are kept, and their log probabilities are converted to “confidences”, which are also stored. Since not every hypothesis is equally good, the confidences can be used to weight the hypotheses during inference. We use an “acoustic-only” approach, in the sense that we do not use any sort of language model or pronunciation dictionary to prune the N -best list.

2.3 Inference

Algorithm 2 describes how the wakeword is detected after the wakeword model has been learned. A voice activity detector (VAD) is used to determine which frames contain speech audio; only these frames are sent to the label model. The VAD thus reduces power consumption by reducing the amount of computation performed by the label model. After the label posteriors are computed by the network, the log probability of each hypothesis in the wakeword model is computed. The CTCForward function returns the log probability of a hypothetical label sequence given the audio by efficiently summing over all possible alignments of the label sequence to the audio [3]. The log probabilities are weighted by their respective confidences before they are summed to obtain a score. If the score is above a certain pre-determined threshold, the wakeword is detected.

For clarity, we have written Algorithm 2 as though the posteriors are only computed after a complete audio \mathbf{x}^{test} has been acquired; it is preferable to reduce latency by computing the posteriors and

Algorithm 1 Learning

Require: $\mathbf{x}^{train,1}, \mathbf{x}^{train,2}, \mathbf{x}^{train,3}, \phi$

- 1: $wake_model := \emptyset$
- 2: **for** i in $(1, 2, 3)$ **do**
- 3: $\pi^{train,i} := f_\phi(\mathbf{x}^{train,i})$
- 4: $beam, beam_scores := CTCBeamSearch(\pi^{train,i}, B)$
- 5: **for** $j = 1$ to N **do**
- 6: $\hat{\mathbf{y}}_j^{train,i} := beam(j)$
- 7: $\log p_\phi(\hat{\mathbf{y}}_j^{train,i} | \mathbf{x}^{train,i}) := beam_scores(j)$
- 8: $w_j^{train,i} := -\frac{1}{\log p_\phi(\hat{\mathbf{y}}_j^{train,i} | \mathbf{x}^{train,i})}$
- 9: $wake_model := wake_model \cup (\hat{\mathbf{y}}_j^{train,i}, w_j^{train,i})$
- 10: **end for**
- 11: **end for**
- 12: **return** $wake_model$

updating the hidden states as each speech frame becomes available from the VAD. Likewise, the forward algorithm can ingest a slice of π^{test} at each timestep to compute that timestep’s forward probabilities.

Algorithm 2 Inference

Require: $\mathbf{x}^{test}, wake_model, \phi$

- 1: $\pi^{test} := f_\phi(\mathbf{x}^{test})$
- 2: $score := 0$
- 3: **for** $(\hat{\mathbf{y}}, w)$ in $wake_model$ **do**
- 4: $\log p_\phi(\hat{\mathbf{y}} | \mathbf{x}^{test}) := CTCForward(\pi^{test}, \hat{\mathbf{y}})$
- 5: $score := score + \log p_\phi(\hat{\mathbf{y}} | \mathbf{x}^{test}) \cdot w$
- 6: **end for**
- 7: **return** $score$

2.4 Runtime requirements

DONUT is fast and suitable for running online on an embedded device. The memory, storage, and computational requirements of running DONUT online can be broken down into two parts: running the label model and running the wakeword model.

The runtime requirements are dominated by the label model (the neural network). The complexity of running the neural network is $O(nT)$, where n is the number of parameters and T is the duration of the audio in frames. We use an RNN with frame stacking [18]: that is, pairs of contiguous acoustic frames are stacked together so that the RNN operates at 50 Hz instead of 100 Hz, cutting the number of operations in half at the expense of slightly more input-hidden parameters in the first layer.

The wakeword model requires little storage, as it consists of just $3N$ short strings and one real-valued confidence for each string. The CTC forward algorithm requires $O(UT)$ operations to process a single label sequence. If the algorithm is run separately for N hypotheses, and the hypotheses have length U on average, then $O(NUT)$ operations are required. The number of operations could be reduced by identifying and avoiding recomputing shared terms for the forward probabilities (e.g. using a lattice [19]), at the cost of a more complicated implementation. However, since N and U are small values, this kind of optimization is not crucial. (In the experiments described below, n is 168k, N is 10, and U is on average 10, so it is apparent that in general $O(nT) \gg O(NUT)$.) The system requires $O(NU)$ memory to store the forward probabilities for a single timestep; the memory for the previous timestep can be overwritten with the current timestep after the current forward probabilities have been computed.

Truth: "of dress" (AH V D R EH S)

$\hat{y}_{1,\dots,N}^{train,1}, w_{1,\dots,N}^{train,1}$:
 AH N D R AE S: 0.274
 AH N D AE S: 0.252
 AH N D ER AE S: 0.238
 ...

$\hat{y}_{1,\dots,N}^{train,2}, w_{1,\dots,N}^{train,2}$:
 AH N D ER S: 0.310
 AH N D S: 0.258
 AH N D EH S: 0.249
 ...

$\hat{y}_{1,\dots,N}^{train,3}, w_{1,\dots,N}^{train,3}$:
 AH V D AE S T: 0.279
 AH V D ER AE S T: 0.260
 AH V D R AE S T: 0.256
 ...

Test phrase	Class	Score
of distress	non-wakeword (confusing)	-136.56
of dressing	non-wakeword (confusing)	-87.74
of these	non-wakeword (confusing)	-142.07
or less	non-wakeword (confusing)	-207.12
others	non-wakeword (non-confusing)	-179.99
into a fever	non-wakeword (non-confusing)	-221.66
torn by	non-wakeword (non-confusing)	-229.12
to be captured your	non-wakeword (non-confusing)	-270.93
of dress	wakeword	-98.43
of dress	wakeword	-72.45
of dress	wakeword	-114.33
of dress	wakeword	-92.05
of dress	wakeword	-69.63
of dress	wakeword	-120.63
of dress	wakeword	-102.24
of dress	wakeword	-72.21

Figure 2: Example of a wakeword model generated from three examples of the target “of dress” (left) and an episode for that target (right). All examples in this episode are classified correctly given a detection threshold of -121, except for “of dressing”, which contains the wakeword as a substring.

3 Experiments

3.1 Data

All audio data in our experiments is sampled at 16,000 Hz and converted to sequences of 41-dimensional Mel filterbank (FBANK) feature vectors using a 25 ms window with a stride of 10 ms. Here, we describe the two types of datasets used in our experiments: the dataset used to train the label models, and the datasets used to train and test the wakeword detectors.

Label dataset We used LibriSpeech [20], an English large vocabulary continuous speech recognition (LVCSR) dataset, to train label models. We used the Montreal Forced Aligner [21] to obtain phoneme-level transcripts written in ARPAbet of the 100- and 360-hour subsets of the dataset. We trained a unidirectional GRU network with 3 layers and 96 hidden units per layer (168k parameters) on LibriSpeech with CTC using the phoneme-level transcripts.

Wakeword datasets We created two wakeword datasets: one based on the 500-hour subset of LibriSpeech (LibriSpeech-Fewshot) and one based on crowdsourced English recordings (English-Fewshot). Both datasets are composed of a number of few-shot learning “episodes”. Each episode contains support examples and test examples. The support set contains three examples of the target phrase spoken by a single speaker. The test set contains a number of positive and negative examples. An example of an episode is shown in Fig. 2. The episodes are split into one subset for hyperparameter tuning and another subset for reporting performance.

To create the LibriSpeech-Fewshot dataset, we split the LibriSpeech recordings into short phrases between 500 ms and 1,500 ms long, containing between one and four words. These short phrases were selected and grouped together to form 6,047 episodes. The test set contains eight positive examples by the same speaker and 24 negative examples by random speakers. Of the negative examples, twenty are phonetically similar (“confusing”), and four are phonetically dissimilar (“non-confusing”). To produce the confusing examples, we generated a phoneme-level transcript for each example, calculated the phoneme edit distance between the target phrase and all other available phrases, and chose the 20 phrases with the lowest phoneme edit distance. The non-confusing examples were chosen at random from the remaining phrases.

To create the English-Fewshot dataset, we used crowdsourcing to record speakers saying phrases consisting of “Hello” followed by another word: for example, “Hello Computer”. Like the LibriSpeech-

Table 1: Performance on English-Fewshot of query-by-example methods.

Method	Same speaker (confusing)		Same speaker (non-confusing)		Different speaker (non-confusing)	
	EER	AUC	EER	AUC	EER	AUC
DTW (FBANK)	24.2%	0.839	32.4%	0.745	19.3%	0.893
DTW (posteriorgram)	20.8%	0.877	17.3%	0.912	11.0%	0.956
DONUT	7.8%	0.975	7.3%	0.977	3.7%	0.993

Fewshot dataset, this dataset has positive examples from the same speaker and negative examples from different speakers; however, here there are also negative examples from the same speaker, so as to show that the models are not simply performing speaker verification. Due to data-gathering constraints, we were unable to obtain “imposter” examples in which a different speaker says the target phrase, but we plan to explore this in the future.

All wakeword models used beam width $B = 100$ and kept $N = 10$ hypotheses per training example. We use the receiver operating characteristic (ROC) curve to measure the performance of a wakeword detector. A single detection threshold is used across all episodes. Two performance metrics are reported: the equal error rate (EER; lower is better) and the area-under-ROC-curve (AUC; higher is better) metric. An EER of 0% or an AUC of 1 indicates a perfect classifier.

3.2 Comparison with other query-by-example methods

In the first experiment, we compare the performance of DONUT with two other query-by-example keyword spotting methods: dynamic time warping (DTW) based on the raw FBANK input and DTW based on the posteriorgram (the output of the label model). We used the ℓ_2 norm to compare FBANK features, and we used the distance-like metric suggested in [8] to compare posteriorgram features:

$$d(p, q) = -\log((\lambda u + (1 - \lambda)p) \cdot (\lambda u + (1 - \lambda)q)), \quad (1)$$

where λ is a small positive number (we used $1e-5$) and u is a uniform distribution (a vector with entries equal to $\frac{1}{1+|A|}$; used to prevent $\log(0)$ by smoothing the peaky output distribution). We also tried removing the softmax, using the ℓ_2 norm as the distance metric, and using a label model trained using the framewise cross-entropy loss instead of the CTC loss. None of these modifications improved performance; we report the best result with the CTC model here.

Table 1 shows the performance of the query-by-example methods on English-Fewshot. We report the performance for three separate cases, in decreasing order of difficulty: the cases when the negative examples are 1) confusing and taken from the same speaker, 2) non-confusing and taken from the same speaker, and 3) non-confusing and taken from different speakers. DONUT outperforms both DTW methods in all three cases.

3.3 Comparison with query-by-string method

In this experiment, we compare the performance of our method with the performance of conventional CTC keyword spotting when the “true” label sequence is provided (e.g., by the user through a text interface). The phoneme sequence for each phrase in the LibriSpeech-Fewshot dataset was obtained using forced alignment and used as the wakeword model for each episode. Table 2 shows that for phonetically confusing examples, DONUT outperforms the text-based approach, and for non-confusing examples, the two approaches perform roughly the same, with the text-based approach performing very slightly better. This result indicates that not only does DONUT provide a more convenient interface than query-by-string keyword spotting, it also has the same or even better performance.

3.4 Interpretability of the model

Like conventional CTC keyword spotting, DONUT is interpretable, which makes it easy for a system designer to identify problems with the model and improve it. For example, Fig. 2 shows an example of a wakeword model learned for the phrase “of dress”. In the first two training examples, the network

Table 2: Performance on LibriSpeech-Fewshot compared to query-by-string CTC.

Method	Confusing		Non-confusing	
	EER	AUC	EER	AUC
CTC (query-by-string)	26.9%	0.810	9.6%	0.968
DONUT	21.0%	0.872	9.6%	0.966

Table 3: Impact of number of hypotheses on performance.

Beam width B	# of kept hypotheses N	EER
1 (greedy)	1	4.2%
100	1	4.4%
100	2	4.2%
100	5	3.8%
100	10	3.7%
100	20	3.4%
100	50	3.3%
100	100	3.1%

hears an “N” sound where one would expect the “V” phoneme in “of”. This information can be used to improve the model: one could retrain the label model with more examples short words such as “of” and “on”, to help the model distinguish short sounds more easily. Alternately, it could become apparent after listening to the training examples that for the speaker’s personal accent the phrase does indeed contain an “N” sound.

Debugging the inference phase is also made easier by the use of CTC. It is possible to decode phoneme sequences from the test audio using a beam search, although this is not necessary to do during inference. One could inspect the decoded sequences from an audio that causes a false accept to identify hypotheses that should be removed from the model to make the false accept less likely to occur. If a false reject occurs, one could check whether the wakeword model hypotheses are found in the decoded sequences or if the network hears something completely different.

3.5 Impact of hyperparameters on performance

DONUT has a few hyperparameters: the beam width B , the number of hypotheses kept from the beam search N , the label model ϕ , and the way in which the hypothesis scores are aggregated. Here we explore the impact of these hyperparameters on performance using the English-Fewshot dataset.

Increasing the number of hypotheses generally improves performance (Table 3), though we have found that this may yield diminishing returns. Even a simple greedy search ($B = 1, N = 1$), which can be implemented by picking the top output at each timestep, works fairly well for our system.

With respect to the impact of the choice of label model, we find that label models with lower phoneme error rate (edit distance between the true label sequence and the model’s prediction) for the original corpus they were trained on have a lower error rate for wakeword detection (Table 4). This suggests that making an improvement to the label model can be expected to translate directly to a decrease in EER/increase in AUC.

In the inference algorithm described above (Algorithm 2), the hypotheses’ scores are aggregated by taking a weighted sum, where each weight is the inverse of the log probability of that hypothesis

Table 4: Impact of label model quality on performance.

Model size	Phoneme error rate	EER
2x128 (186k params)	17.5%	5.7%
3x96 (168k params)	15.8%	3.7%
3x512 (4m params)	11.5%	1.4%

given its corresponding training example. Without the weighting, performance was hurt because some hypotheses are a worse fit to the data than others. A more principled approach to aggregating the scores might be to treat the hypotheses' log probabilities from training as log priors and add them to the scores, since multiplying by a prior is equivalent to adding a log prior, and to take the `logsumexp()` of the scores plus their log priors, since adding two probabilities is equivalent to taking the `logsumexp()` of two log probabilities. However, we have found that this does not work as well as the weighted sum approach, perhaps because the `logsumexp()` function acts like `max()` and tends to pick out a single hypothesis instead of smoothly blending the hypotheses.

4 Conclusion

In this paper, we proposed DONUT, an efficient algorithm for online query-by-example keyword spotting using CTC. The algorithm learns a list of hypothetical label sequences from the user's speech during enrollment and uses these hypotheses to score audios at test time. We showed that the model is interpretable, and thus easy to inspect, debug, and tweak, yet at the same time has high accuracy. Because training a wakeword model amounts to a simple beam search, it is possible to train a model on the user's device without uploading a user's private voice data to the cloud.

Our technique is in principle applicable to any domain in which a user would like to teach a system to recognize a sequence of events, such as a melody (a sequence of musical notes) or a gesture (a sequence of hand movements). It would be interesting to see how well the proposed technique transfers to these other domains.

References

- [1] Guoguo Chen, Carolina Parada, and Georg Heigold, "Small-footprint keyword spotting using deep neural networks," *ICASSP*, 2014.
- [2] Samuel Myer and Vikrant Singh Tomar, "Efficient keyword spotting using time delay neural networks," *Interspeech*, 2018.
- [3] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, "Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks," *ICML*, 2006.
- [4] Kyuyeon Hwang, Minjae Lee, and Wonyong Sung, "Online keyword spotting with a character-level recurrent neural network," *arXiv preprint arXiv:1512.08903*, 2015.
- [5] Chris Lengerich and Awni Hannun, "An end-to-end architecture for keyword spotting and voice activity detection," *NIPS*, 2016.
- [6] Yimeng Zhuang, Xuankai Chang, Yanmin Qian, and Kai Yu, "Unrestricted vocabulary keyword spotting using LSTM-CTC," *Interspeech*, 2016.
- [7] Thibault Gisselbrecht, "Machine Learning on Voice: a gentle introduction with Snips Personal Wake Word Detector," 2018, <https://medium.com/snips-ai/machine-learning-on-voice-a-gentle-introduction-with-snips-personal-wake-word-detector-133bd6fb568e>.
- [8] Timothy J Hazen, Wade Shen, and Christopher White, "Query-by-example spoken term detection using phonetic posteriorgram templates," *ASRU*, 2009.
- [9] Yaodong Zhang and James R. Glass, "Unsupervised spoken keyword spotting via segmental DTW on Gaussian posteriorgrams," *ASRU*, 2009.
- [10] Luis J. Rodriguez-Fuentes, Amparo Varona, Mikel Penagarikano, Germán Bordel, and Mireia Diez, "High-performance Query-by-Example Spoken Term Detection on the SWS 2013 evaluation," *ICASSP*, 2014.
- [11] Guoguo Chen, Carolina Parada, and Tara N. Sainath, "Query-by-example keyword spotting using Long Short Term Memory Networks," *ICASSP*, 2015.

- [12] Shane Settle and Karen Livescu, “Discriminative Acoustic Word Embeddings: Recurrent Neural Network-Based Approaches,” *SLT*, 2016.
- [13] Shane Settle, Keith Levin, Herman Kamper, and Karen Livescu, “Query-by-example search with discriminative neural acoustic word embeddings,” *Interspeech*, 2017.
- [14] Alec Radford, Rafal Jozefowicz, and Ilya Sutskever, “Learning to generate reviews and discovering sentiment,” *arXiv preprint arXiv:1704.01444*, 2017.
- [15] Lyan Verwimp, Hugo Van hamme, Vincent Renkens, and Patrick Wambacq, “State Gradients for RNN Memory Analysis,” *Interspeech*, 2018.
- [16] Théodore Bluche, Hermann Ney, Jérôme Louradour, and Christopher Kermorvant, “Framewise and CTC training of neural networks for handwriting recognition,” *ICDAR*, 2015.
- [17] Awni Hannun, “Sequence modeling with CTC,” *Distill*, 2017, <https://distill.pub/2017/ctc>.
- [18] Haşim Sak, Andrew Senior, Kanishka Rao, and Françoise Beaufays, “Fast and Accurate Recurrent Neural Network Acoustic Models for Speech Recognition,” *Interspeech*, 2015.
- [19] Zhehuai Chen, Wei Deng, Tao Xu, and Kai Yu, “Phone synchronous decoding with CTC lattice,” *Interspeech*, 2016.
- [20] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “LibriSpeech: an ASR corpus based on public domain audio books,” *ICASSP*, 2015.
- [21] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, “Montreal Forced Aligner: Trainable text-speech alignment using Kaldi,” *Interspeech*, 2017.