Published in final edited form as:

Proc IEEE Int Conf Acoust Speech Signal Process. 2017 March ; 2017: 2422–2426. doi:10.1109/ICASSP.2017.7952591.

SEQUENCE SEGMENTATION USING JOINT RNN AND STRUCTURED PREDICTION MODELS

Yossi Adi¹, Joseph Keshet¹, Emily Cibelli², and Matthew Goldrick²

¹Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

²Department of Linguistics, Northwestern University, Evanston, IL, USA

Abstract

We describe and analyze a simple and effective algorithm for sequence segmentation applied to speech processing tasks. We propose a neural architecture that is composed of two modules trained jointly: a recurrent neural network (RNN) module and a structured prediction model. The RNN outputs are considered as feature functions to the structured model. The overall model is trained with a structured loss function which can be designed to the given segmentation task. We demonstrate the effectiveness of our method by applying it to two simple tasks commonly used in phonetic studies: word segmentation and voice onset time segmentation. Results suggest the proposed model is superior to previous methods, obtaining state-of-the-art results on the tested datasets.

Index Terms

Sequence segmentation; recurrent neural networks (RNNs); structured prediction; word segmentation; voice onset time

1. INTRODUCTION

Sequence segmentation is an important task for many speech and audio applications such as speaker diarization, laboratory phonology research, speech synthesis, and automatic speech recognition (ASR). Segmentation models can be used as a pre-process step to clean the data (e.g., removing non-speech regions such as music or noise to reduce ASR error [1, 2]). They can also be used as tools in clinically- or theoretically-focused phonetic studies that utilize acoustic properties as a dependent measure. For example, voice onset time, a key feature distinguishing voiced and voiceless consonants across languages [3], is important both in ASR [4], clinical [5], and theoretical studies [6].

Previous work on speech sequence segmentation focuses on generative models such as hidden Markov models (see for example [7] and the references therein); on discriminative methods [2, 8, 9]; or on deep learning [10, 11].

Inspired by the recent work on combined deep network and structured prediction models [12, 13, 14, 15, 16], we would like to further improve performance on speech sequence

segmentation and propose a new efficient joint deep network and structure prediction model. Specifically, we jointly optimize RNN and structured loss parameters by using RNN outputs as feature functions for a structured prediction model. First, an RNN encodes the entire speech utterance and outputs new representation for each of the frames. Then, an efficient search is applied over all possible segments so that the most probable one can be selected. We evaluate this approach using two tasks: word segmentation and voice onset time segmentation. In both tasks the input is a speech segment and the goal is to determine the boundaries of the defined event. We show that the proposed approach outperforms previous methods on these two segmentation tasks.

2. PROBLEM SETTING

In the problem of speech segmentation we are provided with a speech utterance, denoted as $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$, represented as a sequence of acoustic feature vectors, where each $\mathbf{x}_t \in \mathbb{R}^D$ ($1 \leq t \leq T$) is a D -dimensional vector. The length of the speech utterance, T , is not a fixed value, since the input utterances can have different durations.

Each input utterance is associated with a timing sequence, denoted by $\bar{\mathbf{y}} = (y_1, \dots, y_p)$, where p can vary across different inputs. Each element $y_i \in \mathcal{Y}$, where $\mathcal{Y} = \{1, \dots, T\}$ indicates the start time of a new event in the speech signal. We annotate all the possible timing sequence of size p by $\mathcal{Y}^p$.

For example, in *word segmentation* the goal is to segment a word from silence and noise in the signal. In this case the size of $\bar{\mathbf{y}}$ is 2, namely word onset and offset. However, in *phoneme segmentation* the goal is to segment every phoneme in a spoken word. In this case the size of $\bar{\mathbf{y}}$ is different for each input sequence.

Generally, our method is suitable for different sequence size $|\bar{\mathbf{y}}|$. In this paper we focused on $|\bar{\mathbf{y}}| = 2$, and leave the problem of $|\bar{\mathbf{y}}| > 2$ to future work.

3. MODEL DESCRIPTION

We now describe our model in greater detail. First, we present the structured prediction framework and then discuss how it is combined with an RNN.

3.1. Structured Prediction

We consider the following prediction rule with $\mathbf{w} \in \mathbb{R}^d$, such that $\bar{\mathbf{y}}'_w$ is a good approximation to the true label of $\bar{\mathbf{x}}$, as follows:

$$\bar{\mathbf{y}}'_w(\bar{\mathbf{x}}) = \underset{\bar{\mathbf{y}} \in \mathcal{Y}}{\operatorname{argmax}} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \quad (1)$$

Following the structured prediction framework, we assume there exists some unknown probability distribution ρ over pairs $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ where $\bar{\mathbf{y}}$ is the desired output (or reference output) for input $\bar{\mathbf{x}}$. Both $\bar{\mathbf{x}}$ and $\bar{\mathbf{y}}$ are usually structured objects such as sequences, trees, etc. Our goal is to set $\mathbf{w}$ so as to minimize the expected cost, or the *risk*,

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \mathbb{E}_{(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \sim \rho} [\ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}'_w(\bar{\mathbf{x}}))]. \quad (2)$$

This objective function is hard to minimize directly since the distribution ρ is unknown. We use a training set $\mathcal{S} = \{(\bar{\mathbf{x}}_1, \bar{\mathbf{y}}_1), \dots, (\bar{\mathbf{x}}_m, \bar{\mathbf{y}}_m)\}$ of m examples that are drawn i.i.d. from ρ , and replace the expectation in (2) with a mean over the training set.

The cost is often a combinatorial non-convex quantity, which is hard to minimize. Hence, instead of minimizing the cost directly, we minimize a slightly different function called a *surrogate loss*, denoted $\bar{\ell}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}})$, and closely related to the cost. Overall, the objective function in (2) transforms into the following objective function, denoted as F :

$$F(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) = \frac{1}{m} \sum_{i=1}^m \bar{\ell}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) \quad (3)$$

In this work the surrogate loss function is the structural hinge loss [17] defined as

$$\bar{\ell}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) = \max_{\bar{\mathbf{y}}' \in \mathcal{Y}} [\ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}') - \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) + \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}')]$$

Usually, $\phi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is manually chosen using data analysis techniques and involves manipulation on local and global features. In the next subsection we describe how to use an RNN as feature functions.

3.2. Recurrent Neural Networks as Feature Functions

RNN is a deep network architecture that can model the behavior of dynamic temporal sequences using an internal state which can be thought of as memory [18, 19]. RNN provides the ability to predict the current frame label based on the previous frames. Bidirectional RNN is a model composed of two RNNs: the first is a standard RNN while the second reads the input backwards. Such a model can predict the current frame based on both past and future frames. By using the RNN outputs we can jointly train the structured and network models.

Recall our prediction rule in Eq. (1): notice that $\phi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ can be viewed as $\sum_{i=1}^p \phi'(\bar{\mathbf{x}}, y_i)$ where each ϕ can be extracted using different techniques, e.g., hand-crafted, feed-forward neural network, RNNs, etc. We can formulate the prediction rule as follows:

$$\begin{aligned}
\bar{\mathbf{y}}'_w(\bar{\mathbf{x}}) &= \arg\max_{\bar{\mathbf{y}} \in \mathcal{Y}^p} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \\
&= \arg\max_{\bar{\mathbf{y}} \in \mathcal{Y}^p} \mathbf{w}^\top \sum_{i=1}^p \phi'(\bar{\mathbf{x}}, y_i) \\
&= \arg\max_{\bar{\mathbf{y}} \in \mathcal{Y}^p} \mathbf{w}^\top \sum_{i=1}^p \text{RNN}(\bar{\mathbf{x}}, y_i),
\end{aligned} \tag{4}$$

where the RNN can be of any type and architecture. For example, we can use bidirectional RNN and consider ϕ as the concatenation of both outputs $\text{BI-RNN}_{\text{forward}} \oplus \text{BI-RNN}_{\text{backward}}$. This is depicted in Figure 1. We call our model *DeepSegmentor*.

Our goal is to find the model parameters so as to minimize the risk as in Eq. (2). Recall, we use the structural hinge loss function, and since both the loss function and the RNN are differentiable we can optimize them using gradient based methods such as stochastic gradient descent (SGD). In order to optimize the network parameters using the back-propagation algorithm [20], we must find the outer derivative of each layer with respect to the model parameters and inputs.

The derivative of the loss layer with respect to the layer parameters $\mathbf{w}$ for the training example $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is

$$\frac{\partial F}{\partial \mathbf{w}} = \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}_w^\ell) - \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}),$$

where

$$\bar{\mathbf{y}}_w^\ell = \arg\max_{\bar{\mathbf{y}}' \in \mathcal{Y}^p} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}') + \ell(\bar{\mathbf{y}}, \bar{\mathbf{y}}'). \tag{5}$$

Similarly, the derivatives with respect to the layer's inputs are

$$\frac{\partial F}{\partial \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}})} = -\mathbf{w} \quad \frac{\partial F}{\partial \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}')} = \mathbf{w}.$$

The derivatives of the rest of the layers are the same as an RNN model.

4. EXPERIMENTAL RESULTS

We investigate two segmentation problems; word segmentation and voice onset time segmentation. We describe each of them in details in the following subsections.¹

4.1. Word Segmentation

In the problem of word segmentation we are provided with a speech utterance which contains a single word; our goal is to predict its start and end times. The ability to determine these timings is crucial to phonetic studies that measure speaker properties (e.g. response time [23]) or as a preprocessing step for other phonetic analysis tools [11, 10, 9, 8, 24].

4.1.1. Dataset—Our dataset comes from a laboratory study by Fink and Goldrick [23]. Native English speakers were shown a set of 90 pictures. Some participants produced the name of the picture (e.g., saying “cat”, “chair”) while others performed a semantic classification task (e.g., saying “natural”, “man-made”). Productions other than the intended response or disfluencies were excluded. Recordings were randomly assigned to two transcribers who annotated the onset and offset of each word. We analyze a subset of the recordings, including data from 60 participants, evenly distributed across tasks.

4.1.2. Results—We compare our model to an RNN that was trained using the Negative-Log-Likelihood (NLL). The NLL model makes a binary decision in every frame to predict whether there is voice activity or not. Recall, our goal is to find the start and end times of the word; in this task, the RNN leaves us with a distribution over all possible onsets. To account for this, we apply a smoothing algorithm and find the most probable pair of timings.

We trained the DeepSegmentor model using the structured loss function as in (6), denoted as Combined Duration (CD) loss. The motivation for using this function is due to disparities in the manual annotations, which are common and depend both on human errors and objective difficulties in placing the boundaries. Hence we chose a loss function that takes into account the variations in the annotations.

$$\gamma(\bar{y}, \bar{y}') = [|y_1 - y'_1| - \tau]_+ + [|y_2 - y'_2| - \tau]_+, \quad (6)$$

where $[\pi]_+ = \max\{0, \pi\}$, and τ is a user defined tolerance parameter.

We use two layers of bidirectional LSTMs for the DeepSegmentor model with dropout [25] after each recurrent layer. We extracted the 13 Mel-Frequency Cepstrum Coefficients (MFCCs), without the deltas, every 10 ms, and use them as inputs to the network. We optimize the networks using AdaGrad [26]. All parameters were tuned on a dedicated development set for both of the models. As for the NLL models, we trained 4 different models; LSTM with one and two layers, and bidirectional LSTM with one and two layers, denoted as RNN, 2RNN, BI-RNN and BI-2-RNN, respectively. Table 1 summarizes the results for both models.

Besides being efficient and more elegant, DeepSegmentor is superior to the NLL models when measuring (6), with the exception of BI-2-RNN, which was slightly better for the offset measurement.

¹All models were implemented using Torch7 toolkit [21, 22]

4.2. VOT Segmentation

Voice onset time (VOT) is the time between the onset of a stop burst and the onset of voicing. As noted in the introduction, it is widely used in theoretical and clinical studies as well as ASR tasks. In this problem the input is a speech utterance containing a single stop consonant, and the output is the VOT onset and offset times.

We compared our model to two other methods for VOT measurement. First is the *AutoVOT* algorithm [9]. This algorithm follows the structured prediction approach of linear classifier with hand-crafted features and feature-functions. The second algorithm is the *DeepVOT* algorithm [11]. This algorithm uses RNNs with NLL as loss function. Hence, it predicts for each frame whether it is related to the VOT or not. Using the RNN predictions, a dynamic programming algorithm is applied to find the best onset and offset times. Our approach combines both of these methods while jointly training RNN with structured loss function.

4.2.1. Datasets—We use two different datasets. The first one, PGWORDS, is from a laboratory study by Paterson and Goldrick [6]. American English monolinguals and Brazilian Portuguese (L1)-English bilinguals (24 participants each) named a set of 144 pictures. Productions other than the intended label as well as those with code-switching or disfluencies were excluded. VOT of remaining words was annotated by one transcriber.

The second dataset, BB, consists of spontaneous speech from the 2008 season of Big Brother UK, a British reality television show [27, 9]. The speech comes from 4 speakers recorded in the “diary room,” an acoustically clean environment. VOTs were manually measured by two transcribers.

4.2.2. Results — PGWORDS—For the PGWORDS dataset we use two layers of bidirectional LSTMs with dropout after each recurrent layer. We use (6) as our loss function. The input features are the same as in [9, 11]; overall we have 63 features per frame. We optimize the networks using AdaGrad optimization. All parameters were tuned on a dedicated development set. Table 2 summarizes the results using the same loss function as in [9]. Results suggests that DeepSegmentor outperforms the AutoVOT model over all tolerance values. However, when comparing to DeepVOT, the picture is mixed. In the lower tolerance values DeepSegmentor is superior to the DeepVOT while for higher values DeepVOT performs better. We believe these results are due to the DeepVOT being less delicate and solving a much coarser problem than the DeepSegmentor; hence, it performs better when considering high tolerance values. We believe the integration between these two systems, (using DeepVOT as pre-training for the DeepSegmentor), will yield more accurate and robust results. We leave this investigation for future work.

4.2.3. Results — BB—For the BB dataset we use two layers of LSTMs with dropout after each recurrent layer. We have experiences with bidirectional LSTMs as well but only forward LSTM performs better on this dataset. We use (6) as our loss function. We use the same features as in [9, 11], overall we have 51 features per frame. We optimize the networks using AdaGrad optimization. All parameters were tuned on a dedicated development set. Table 3 summarize the results using the loss function as in [9]. It is worth notice that we see

the same behavior on this dataset as well, regarding the DeepVOT preforms better then the DeepSegmentor in hight tolerance values.

5. FUTURE WORK

Future work will explore timing sequence of length greater than 2 - for instance, in phoneme segmentation, where the sequence varies across training examples. The model's robustness to noise and length as well as its ability to generalize are also key areas of future development. We would therefore like to explore training the model in two stages: first as a multi-class version and then fine-tuning using structured loss. With respect to machine learning, future directions include the effect of network size, depth, and loss function on model performance.

In this paper we present a new algorithm for speech segmentation and evaluate its performance to two different tasks. The proposed algorithm combines structured loss function with recurrent neural networks and outperforms current state-of-the-art methods.

Acknowledgments

Supported in part by NIH grant 1R21HD077140.

References

1. Kubala, Francis, Anastasakos, Tasos, Jin, Hubert, Nguyen, Long, Schwartz, Richard. Transcribing radio news. ICSLP. 1996; 2:598–601.
2. Rybach, David, Gollan, Christian, Schluter, Ralf, Ney, Hermann. Audio segmentation for speech recognition using segment features. ICASSP. 2009:4197–4200.
3. Lisker L, Abramson A. A cross-language study of voicing in initial stops: acoustical measurements. Word. 1964; 20:384–422.
4. Hansen JHL, Gray SS, Kim W. Automatic voice onset time detection for unvoiced stops (/p/,/t/,/k/) with application to accent classification. Speech Commun. 2010; 52:777–789.
5. Auzou P, Ozsancak C, Morris RJ, Jan M, Eustache F, Hannequin D. Voice onset time in aphasia, apraxia of speech and dysarthria: a review. Clin Linguist Phonet. 2000; 14:131–150.
6. Paterson, Nattalia. PhD thesis. Northwestern University; 2011. Interactions in Bilingual Speech Processing.
7. Toledano, Doroteo Torre, Hernández Gómez, Luis A., Grande, Luis Villarrubia. Automatic phonetic segmentation. IEEE transactions on speech and audio processing. 2003; 11(6):617–625.
8. Keshet, Joseph, Shalev-Shwartz, Shai, Singer, Yoram, Chazan, Dan. A large margin algorithm for speech-to-phoneme and music-to-score alignment. IEEE Transactions on Audio, Speech, and Language Processing. 2007; 15(8):2373–2382.
9. Sonderegger, Morgan, Keshet, Joseph. Automatic measurement of voice onset time using discriminative structured prediction. JASA. 2012; 132(6):3965–3979.
10. Adi, Yossi, Keshet, Joseph, Goldrick, Matthew. Vowel duration measurement using deep neural networks. MLSP. 2015:1–6.
11. Adi, Yossi, Keshet, Joseph, Dmitrieva, Olga, Goldrick, Matt. Automatic measurement of voice onset time and prevoicing using recurrent neural networks.
12. Do, Trinh, Arti, Thierry, et al. Neural conditional random fields. AISTATS. 2010:177–184.
13. Zheng, Shuai, Jayasumana, Sadeep, Romera-Paredes, Bernardino, Vineet, Vibhav, Su, Zhizhong, Du, Dalong, Huang, Chang, Torr, Philip HS. Conditional random fields as recurrent neural networks. ICCV. 2015:1529–1537.

14. Chen, Liang-Chieh, Schwing, Alexander G., Yuille, Alan L., Urtasun, Raquel. Learning deep structured models. ICML. 2015
15. Kiperwasser, Eliyahu, Goldberg, Yoav. Simple and accurate dependency parsing using bidirectional lstm feature representations. 2016 arXiv preprint.
16. Lample, Guillaume, Ballesteros, Miguel, Subramanian, Sandeep, Kawakami, Kazuya, Dyer, Chris. Neural architectures for named entity recognition. 2016 arXiv preprint.
17. Tsochantaridis, Ioannis, Joachims, Thorsten, Hofmann, Thomas, Altun, Yasemin. Large margin methods for structured and interdependent output variables. JMLR. 2005:1453–1484.
18. Elman, Jeffrey L. Distributed representations, simple recurrent networks, and grammatical structure. Machine learning. 1991; 7(2–3):195–225.
19. Graves, Alan, Mohamed, Abdel-rahman, Hinton, Geoffrey. Speech recognition with deep recurrent neural networks. ICASSP. 2013:6645–6649.
20. Rumelhart, David E., Hinton, Geoffrey E., Williams, Ronald J. Learning representations by back-propagating errors. Cognitive modeling. 1988; 5(3):1.
21. Collobert, Ronan, Kavukcuoglu, Koray, Farabet, Clément. Torch7: A matlab-like environment for machine learning. BigLearn, NIPS Workshop; 2011. number EPFL-CONF-192376
22. Léonard, Nicholas, Waghmare, Sagar, Wang, Yang. rnn: Recurrent library for torch. 2015 arXiv preprint.
23. Fink, Angela. PhD thesis. Northwestern University; 2016. The Role of Domain-General Executive Functions, Conceptualization, and Articulation during Spoken Word Production.
24. Rosenfelder, Ingrid, Fruehwald, Josef, Evanini, Keelan, Seyfarth, Scott, Gorman, Kyle, Prichard, Hilary, Yuan, Jiahong. Fave (forced alignment and vowel extraction). 2014 Program suite v1.2.2.
25. Hinton, Geoffrey E., Srivastava, Nitish, Krizhevsky, Alex, Sutskever, Ilya, Salakhutdinov, Ruslan. Improving neural networks by preventing co-adaptation of feature detectors. CoRR. 2012
26. Duchi, John, Hazan, Elad, Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. JMLR. 2011; 12:2121–2159.
27. Bane, Max, Graff, Peter, Sonderegger, Morgan. Longitudinal phonetic variation in a closed system. Proc CLS. 2010; 46:43–58.

Fig. 1.
An illustration for using BI-RNN as feature functions. We search through all possible locations and predict the one with the highest score. In this example the target timing sequence is (1, 6) and the predicted timing sequence is (2, 5).

Table 1

The mean loss for NLL models and DeepSegmentor for the word segmentation task. Results are reported for the onset, offset and overall CD separately. The loss function was measured using (6) (with $\tau=0$) in frames of 10ms.

	RNN	2-RNN	BI-RNN	BI-2-RNN	DeepSeg.
Onset	6.0	5.84	2.88	3.48	2.02
Offset	9.43	8.92	4.46	3.75	3.96
CD	15.42	14.76	7.35	7.24	5.98

Table 2

Proportion of differences between automatic and manual measures falling at or below a given tolerance value (in msec). For example, for DeepVOT, the difference between automatic and manual measurements in the test set was 2 msec or less in 53.8% of examples. These results are for the PGWORDS dataset.

Model	t	2	t	5	t	10	t	15	t	25	t	50
AutoVOT		49.1		81.3		93.9		96.0		97.2		98.1
DeepVOT		53.8		91.6		97.6		98.7		99.6		100
DeepSeg		78.2		94.1		97.1		98.6		99.1		99.4

Table 3

Proportion of differences between automatic and manual measures falling at or below a given tolerance value (in msec). These results are for the BB dataset.

Model	t	2	t	5	t	10	t	15	t	25	t	50
AutoVOT		59.1		80.5		89.9		94.3		96.8		98.1
DeepVOT		60.3		84.2		94.3		94.9		98.1		98.7
DeepSeg.		64.8		85.5		94.3		95.0		96.2		97.5