

# PHONEME BOUNDARY DETECTION USING LEARNABLE SEGMENTAL FEATURES

Felix Kreuk<sup>1</sup>, Yaniv Sheena<sup>2</sup>, Joseph Keshet<sup>1</sup>, Yossi Adi<sup>3</sup>

<sup>1</sup>Bar-Ilan University, Ramat-Gan, Israel

<sup>2</sup>Facebook Inc.

<sup>3</sup>Facebook AI Research

## ABSTRACT

Phoneme boundary detection plays an essential first step for a variety of speech processing applications such as speaker diarization, speech science, keyword spotting, etc. In this work, we propose a neural architecture coupled with a parameterized structured loss function to learn segmental representations for the task of phoneme boundary detection. First, we evaluated our model when the spoken phonemes were not given as input. Results on the TIMIT and Buckeye corpora suggest that the proposed model is superior to the baseline models and reaches state-of-the-art performance in terms of F1 and R-value. We further explore the use of phonetic transcription as additional supervision and show this yields minor improvements in performance but substantially better convergence rates. We additionally evaluate the model on a Hebrew corpus and demonstrate such phonetic supervision can be beneficial in a multi-lingual setting.

**Index Terms**— Sequence segmentation, recurrent neural networks (RNNs), structured prediction, phoneme boundary detection

## 1. INTRODUCTION

*Phoneme Boundary Detection* or *Phoneme Segmentation* plays an essential first step for a variety of speech processing applications such as speaker diarization [1], speech science [2, 3], keyword spotting [4], Automatic Speech Recognition [5, 6], etc.

Such segmentations are often modeled by either supervised or unsupervised methods. In the unsupervised setting, only the speech signal is given as input [7, 8], whereas in the supervised setting, the speech signal is accompanied by the target boundaries as supervision. Under the supervised setting, the set of pronounced or presumed phonemes is often provided additionally as input. This setting is referred to as *forced alignment*. In cases where no phonemes are provided, we often denote the setup as *text-independent* phoneme segmentation [9, 10]. In this study, we focus on the latter, namely phoneme boundary detection with no phonetic features.

Inspired by the work on learning segmental features [11, 12, ?], we suggest learning segmental representations using

neural models for both phoneme boundaries and phoneme segments to detect phoneme boundaries accurately.

Specifically, we jointly optimize a Recurrent Neural Network (RNN) and structured loss parameters by using RNN outputs as feature functions for a structured model. First, the RNN encodes the entire speech utterance and outputs a new representation for each frame. Then, an efficient search is applied over all possible segments so that the most probable one can be selected. To compare with prior work, we evaluate this approach using TIMIT [13] and Buckeye [14] datasets. The proposed approach outperforms the baseline models on both datasets and reaches State-of-the-Art (SOTA) performance. We additionally experiment with leveraging phoneme information as secondary supervision to the model and demonstrate that such an approach can be beneficial for both performance and convergence speed. Lastly, we demonstrated that such supervision could be beneficial when detecting phoneme boundaries from different languages, by training on English and testing on Hebrew.

## 2. RELATED WORK

The problem of phoneme boundary detection was explored in various settings. In the unsupervised setting, also known as *blind phoneme segmentation*, the speech signal is provided by itself with no additional phonemes nor boundaries as supervision. Traditionally, signal processing techniques were used to find spectral changes in the signal to detect phoneme boundaries [8, 15, 16, 17, 18] and the references therein. Recently, the authors of [7] suggested training a next-frame prediction model using either an approximated Markov model or RNN to identify potential phoneme transitions at regions of high error.

Under the supervised setting, the most common approach is the forced alignment setup. Models that follow such an approach involve with Hidden Markov Models or different structured prediction algorithms using handcrafted features [19, 20]. The main drawback of forced alignment methods is the need for phoneme annotations also at inference time. This requirement may limit the applicability of such techniques to a monolingual setting due to mismatched phoneme sets, i.e., it is unclear how to handle unseen

---

Work was conducted while Yaniv was at Bar-Ilan University

phonemes from a foreign language.

Another supervised setting is the text-independent phoneme boundary detection, in which the model is provided with phoneme boundaries as supervision but without information about the uttered phonemes. Most previous works that follow such setup consider the problem as a binary classification task where one label is associated with phoneme boundaries annotations, and another one for the rest of the signal. The authors in [21] used a kernel-based method, composed of six radial basis function support vector machines under such setup, while the authors in [22] suggested using RNN followed by a peak detection algorithm to predict phoneme boundaries. The main drawback of such modeling is that it ignores the internal structure of the output while considering boundaries to be conditionally independent.

Another related lines of work to be noticed is the work of learnable segmental features in ASR [23], dependency parsing [11], and word segmentation [12].

### 3. PROBLEM SETTING

In the problem of phoneme boundary detection 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 input utterances can have different duration.

Each input utterance is associated with a timing sequence, denoted by  $\bar{\mathbf{y}} = (y_1, \dots, y_k)$ , where  $k \leq T$  is the number of segments, and 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 denote by  $\mathcal{Y}^*$  the set of all finite-length sequences over  $\mathcal{Y}$ . Although  $\bar{\mathbf{y}} \in \mathcal{Y}^*$  can be of any length, we assume its length to be bounded by  $T$ .

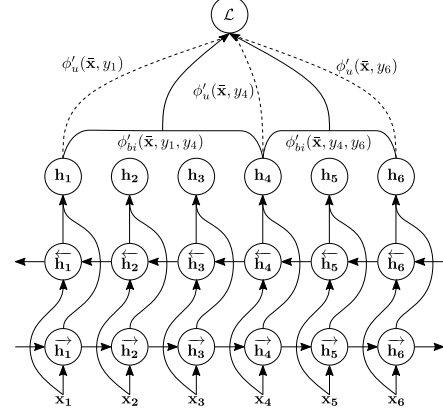
### 4. MODEL DESCRIPTION

Following the structured prediction framework [24], also known as Energy Based Models [25], consider the following prediction rule with  $\mathbf{w} \in \mathbb{R}^d$ , such that  $\bar{\mathbf{y}}_{\mathbf{w}}$  is a good approximation to the true label of  $\bar{\mathbf{x}}$ , as follows:

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

Where  $\phi$  is a mapping function from the set of input objects and target labels to a real vector of length  $d$ . We assume there exists some unknown probability distribution  $\rho$  over pairs  $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$  where  $\bar{\mathbf{y}}$  is the desired output (or reference output) for  $\bar{\mathbf{x}}$ . We are provided with 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  $\rho$ . Our goal is to set  $\mathbf{w}$  so as to minimize the expected empirical cost, or the *empirical risk*,

$$\mathcal{L}(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) = \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) \quad (2)$$



**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 score is calculated for timing sequence is (1, 4, 6).

where we define  $\ell(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}})$  to be the hinge loss function as follows,

$$\ell(\mathbf{w}, \bar{\mathbf{x}}, \bar{\mathbf{y}}) = \max_{\bar{\mathbf{y}}' \in \mathcal{Y}^*} [1 - \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) + \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}')] ]$$

We assume  $\phi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$  can be decomposed as  $\sum_{i=1}^k \phi'(\bar{\mathbf{x}}, y_i)$  where each  $\phi'$  can be extracted using different techniques, e.g., hand-crafted, feed-forward neural network, RNNs, etc. Traditionally,  $\phi'$  was manually chosen using data analysis techniques and involved manipulations of local and global features.

Notice, such decomposition implicitly assumes the output boundaries to be conditionally independent. However, considering the previous boundary prediction provides a lot of information regarding the next boundary. For example, when predicting a vowel offset boundary, considering the vowel onset boundary may provide some insight regarding the typical vowel length. An example of such segmental features might be the mean energy of the spectrum or the average fundamental frequency,  $f_0$ , value in specific ranges [2]. Hence, we adjust the objective as follows,

$$\begin{aligned} \bar{\mathbf{y}}'_{\mathbf{w}}(\bar{\mathbf{x}}) &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^*} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \\ &= \operatorname{argmax}_{\bar{\mathbf{y}} \in \mathcal{Y}^*} \mathbf{w}^\top \left( \sum_{i=1}^k \phi'_u(\bar{\mathbf{x}}, y_i) + \sum_{j=1}^{k-1} \phi'_{bi}(\bar{\mathbf{x}}, y_j, y_{j+1}) \right) \end{aligned} \quad (3)$$

In this work, we consider  $\phi'_u$  and  $\phi'_{bi}$  as *Learnable Segmental Features*. One can explore various types of neural network architectures to model  $\phi'_u$  and  $\phi'_{bi}$ . We set  $\phi'_u(\bar{\mathbf{x}}, y_i)$  to be a bidirectional RNN output at  $y_i$ , and  $\phi'_{bi}(\bar{\mathbf{x}}, y_j, y_{j+1})$  as the sum of the RNN outputs between  $y_j$  and  $y_{j+1}$ . An illustration of the proposed model depicted in Figure 1. Notice

**Table 1.** Comparison of phoneme segmentation models. Precision (P) and recall (R) are calculated with tolerance value of 20 ms

	Model	P	R	F1	R-val
TIMIT	King <i>et al.</i> [21]	87.0	84.8	85.9	87.8
	Franke <i>et al.</i> [22]	91.1	88.1	89.6	90.8
	SEGFEAT	<b>94.03</b>	<b>90.46</b>	<b>92.22</b>	<b>92.79</b>
Buckeye	Franke <i>et al.</i> [22]	<b>87.8</b>	83.3	85.5	87.17
	SEGFEAT	85.4	<b>89.12</b>	<b>87.23</b>	<b>88.76</b>

that the  $\text{argmax}$  operator in Equation 3 may search over exponential number of target segmentations. Luckily, this can be implemented efficiently using a dynamic-programming algorithm, and refer the reader to our implementation for more details. The source code will be available under the following repository: [hyperurl.co/2w206g](https://hyperurl.co/2w206g)

Our goal is to find the model parameters so as to minimize the risk as in Equation 2. Since all components described above are differentiable we can jointly optimize them using gradient based methods such as stochastic gradient descent (SGD). We denote the proposed model as SEGFEAT .

## 5. EXPERIMENTS

In this section, we provide a detailed description of the experiments. We start by presenting the experimental setting. Then we outline the evaluation method. We conclude this section with the boundary detection results and ablation.


### 5.1. Experimental setup

We implemented our model using a 2-layer bidirectional LSTM network. The final score outputs  $\phi'_u$  and  $\phi'_{bi}$  were calculated using a 2-layer fully-connected network that follows the RNN. The proposed model was trained for 150 epochs using Adam optimizer with a learning rate of  $1e^{-4}$ .

We evaluated our model on both TIMIT and Buckeye corpora. For the TIMIT corpus, we used the standard train/test split, where we randomly sampled 10% of the training set for validation. For Buckeye, we split the corpus at the speaker level into training, validation, and test sets with a ratio of 80/10/10. Similarly to [22], we split long sequences into smaller ones by cutting during noises, silences, and untranscribed segments. Overall, each sequence started and ended with a maximum of 20 ms of non-speech<sup>1</sup>. For both corpora, we extracted 13 Mel-Frequency Cepstrum Coefficients (MFCCs), with delta and delta-delta features every 10 ms, with a processing window size of 10ms.

Moreover, we concatenated four additional features based on the spectral changes between adjacent frames, using

<sup>1</sup>All experiments using Buckeye dataset were done at Bar-Ilan university.



**Fig. 2.** F1-score as a function of update steps for different models on the TIMIT validation set.

MFCCs to represent the spectral properties of the frames. Define  $D_{t,j} = d(\mathbf{a}_{t-j}, \mathbf{a}_{t+j})$  to be the Euclidean distance between the MFCC feature vectors  $\mathbf{a}_{t-j}$  and  $\mathbf{a}_{t+j}$ , where  $\mathbf{a}_t \in \mathbb{R}^{39}$  for  $1 \leq t \leq T$ . The features are denoted by  $D_{t,j}$ , for  $j \in \{1, 2, 3, 4\}$ . We observed this set of features greatly improves performance over the standard MFCC features.

### 5.2. Evaluation method

Following previous work on phoneme boundary detection [8, 22, 21, 7], we evaluated the performance of the proposed models and baseline models using precision ( $P$ ), recall ( $R$ ) and F1-score with a tolerance level of 20 ms. A drawback of the F1-score measurement is that a high recall and a low precision might yield relatively a high F1-score (e.g., see Section 4.2 in [7]). The authors in [26] proposed a complementary metric denoted as  $R$ -value:

$$R\text{-value} = 1 - \frac{|r_1| + |r_2|}{2} \quad (4)$$

$$r_1 = \sqrt{(1-R)^2 + (OS)^2}, \quad r_2 = \frac{(-OS + R - 1)}{\sqrt{2}}$$

where  $OS$  is an over-segmentation measure, defined as  $OS = R/P - 1$ . Overall the performance is presented in terms of Precision, Recall, F-score and  $R$ -value.

### 5.3. Results

In Table 1 we compare the proposed model against two baselines: Franke *et al.* [22], and King *et al.* [21]. Results suggest that our proposed model is superior to the baseline models over all metrics on both corpora with one exception of the Precision metric on the Buckeye corpus.

In the next subsection, we conduct an ablation study to explore the effect of using phonetic annotations as additional supervision.

### 5.4. Ablation study

Recall that in the task of phoneme boundary detection, we are often provided with the pronounced unaligned phoneme

**Table 2.** Models performance on TIMIT using different sets of loss function.

Loss	P	R	F1	R-val
BIN	91.1	88.1	89.6	90.8
BIN + PHN	<b>96.6</b>	85.0	90.04	89.33
SEGFEAT	94.03	90.46	92.22	92.79
SEGFEAT +PHN	92.98	92.33	92.66	93.69
SEGFEAT +PHN +BIN	92.67	<b>93.03</b>	<b>92.85</b>	<b>93.91</b>

**Table 3.** Comparison of the proposed model against forced-alignment algorithms.

Model	P	R	F1	R-val
McAuliffe (unsup.) [20]	83.9	81.6	82.7	85.16
Keshet (sup.) [19]	90	82.2	85.9	79.51
SEGFEAT	<b>94.03</b>	<b>90.46</b>	<b>92.22</b>	<b>92.79</b>

sequence. For example, for the timing sequence  $\bar{y} = (19, 23, 40, 45)$ , we additionally get the phoneme sequence  $\bar{p} = (/sil/, /b/, /ae/, /d/, /sil/)$ . As opposed to the forced alignment setup where phonemes are used as additional inputs to the model, one could leverage the provided phonemes as supervision. Such annotations should be provided during the training phase only; hence, they do not restrict the model by any phonetic features. To explore the effect of using phoneme annotations as another supervision we augment the original training objective in Eq. 2 with a phoneme classification loss (Negative Log-likelihood) at each time frame, and denote this loss function by PHN. For completeness, we additionally conducted experiments with including a binary boundary detection loss as suggested in [22, 21], and denote it by BIN.

We trained several models while optimizing different sets of loss functions on the TIMIT corpus. Results are summarized in Table 2. Such comparison sheds light on the effect each loss function has on model performance.

### 5.5. Forced alignment comparison

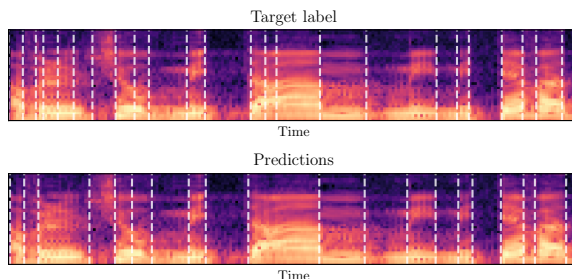
Lastly, in this subsection we compare our model against SOTA forced-alignment models [19, 20]. Recall, forced-alignment is defined as the task of finding the phoneme boundaries of a given set of phonemes, this task may be supervised or unsupervised. Notice that while these models are provided with the pronounced phoneme sequence at test time, the only input to our model is the speech signal. All models were trained and tested on TIMIT dataset. Results are summarized in Table 3.

Leveraging phoneme supervision under the binary boundary detection setting, i.e., optimizing both PHN and BIN loss functions, reaches the best Precision score. This, however, comes at the cost of the worst Recall and R-value. Contrarily, leveraging such supervision using learnable segmental features yields the best Recall, F1-score, and R-value.

While all SEGFEAT models reached a comparable F1-score on the test set, incorporating phoneme supervision im-

**Table 4.** An ablation study on the effect of the PHN loss on Hebrew language.

Model	P	R	F1	R-val
SEGFEAT w/o PHN Loss	<b>83.58</b>	79.2	81.24	83.67
SEGFEAT w PHN Loss	83.11	<b>81.66</b>	<b>82.38</b>	<b>84.92</b>



**Fig. 3.** Example of segmentation result on an Hebrew utterance using an English trained model.

proves R-value and convergence speed by  $\sim 30\%$ . Additionally, adding the BIN loss function to the optimization had a minor effect on the model performance. Figure 2 depicts F1-score as a function of model iterations for different models.

### 5.6. Cross-language phoneme boundary detection

One concern when using the phoneme classification loss is convergence to a language-dependent model. To further explore that we evaluated the performance of two models that were trained on American English (TIMIT) and tested on a Hebrew language corpus [27]. The Hebrew language has only partial overlap in its phonetic set to English, and the utterances in the Hebrew corpus consist of both read and spontaneous speech. Results are summarized in Table 4.

Interestingly, results suggest that incorporating the phoneme classification loss with segmental features slightly improved overall performance in comparison to using segmental features alone. An example of the model predictions together with the target label is depicted in Figure 3.

## 6. DISCUSSION AND FUTURE WORK

We suggest using learnable segmental features for the task of phoneme boundary detection. We showed that such a phoneme-independent approach is superior to the baseline methods and can reach SOTA results on the TIMIT and Buckeye corpora. Moreover, we demonstrated that such a text-independent model could be efficiently applied to the Hebrew language, while it was trained on American English data.

For future work, we would like to explore the cross-domain setting further and while providing a systematic comparison between different languages. Additionally, we would like to examine the usability of the model for low-resource settings.

## 7. REFERENCES

- [1] Mohammad H Moattar and Mohammad M Homayounpour, "A review on speaker diarization systems and approaches," *Speech Communication*, vol. 54, no. 10, pp. 1065–1103, 2012.
- [2] Yossi Adi, Joseph Keshet, Emily Cibelli, Erin Gustafson, Cynthia Clopper, and Matthew Goldrick, "Automatic measurement of vowel duration via structured prediction," *The Journal of the Acoustical Society of America*, vol. 140, no. 6, pp. 4517–4527, 2016.
- [3] Yossi Adi, Joseph Keshet, and Matthew Goldrick, "Vowel duration measurement using deep neural networks," in *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2015, pp. 1–6.
- [4] Joseph Keshet, David Grangier, and Samy Bengio, "Discriminative keyword spotting," *Speech Communication*, vol. 51, no. 4, pp. 317–329, 2009.
- [5] Francis Kubala, Tasos Anastasakos, Hubert Jin, Long Nguyen, and Richard Schwartz, "Transcribing radio news," in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP'96*. IEEE, 1996, vol. 2, pp. 598–601.
- [6] David Rybach, Christian Gollan, Ralf Schluter, and Hermann Ney, "Audio segmentation for speech recognition using segment features," in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2009, pp. 4197–4200.
- [7] Paul Michel, Okko Räsänen, Roland Thiolliere, and Emmanuel Dupoux, "Blind phoneme segmentation with temporal prediction errors," *arXiv preprint arXiv:1608.00508*, 2016.
- [8] Sorin Dusan and Lawrence Rabiner, "On the relation between maximum spectral transition positions and phone boundaries," in *Ninth International Conference on Spoken Language Processing*, 2006.
- [9] Guido Aversano, Anna Esposito, and M Marinaro, "A new text-independent method for phoneme segmentation," in *Proceedings of the 44th IEEE 2001 Midwest Symposium on Circuits and Systems. MWSCAS 2001 (Cat. No. 01CH37257)*. IEEE, 2001, vol. 2, pp. 516–519.
- [10] Lijiang Chen, Xia Mao, and Hong Yan, "Text-independent phoneme segmentation combining egg and speech data," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 24, no. 6, pp. 1029–1037, 2016.
- [11] Eliyahu Kiperwasser and Yoav Goldberg, "Simple and accurate dependency parsing using bidirectional lstm feature representations," *Transactions of the Association for Computational Linguistics*, vol. 4, pp. 313–327, 2016.
- [12] Yossi Adi, Joseph Keshet, Emily Cibelli, and Matthew Goldrick, "Sequence segmentation using joint rnn and structured prediction models," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2422–2426.
- [13] J. S Garofolo, "Timit acoustic phonetic continuous speech corpus," *Linguistic Data Consortium, 1993*, 1993.
- [14] M. A Pitt et al., "The buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability," *Speech Communication*, vol. 45, no. 1, pp. 89–95, 2005.
- [15] Yago Pereiro Estevan, Vincent Wan, and Odette Scharenborg, "Finding maximum margin segments in speech," in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*. IEEE, 2007, vol. 4, pp. IV–937.
- [16] Dac-Thang Hoang and Hsiao-Chuan Wang, "Blind phone segmentation based on spectral change detection using legendre polynomial approximation," *The Journal of the Acoustical Society of America*, vol. 137, no. 2, pp. 797–805, 2015.
- [17] George Almpandis and Constantine Kotropoulos, "Phonemic segmentation using the generalised gamma distribution and small sample bayesian information criterion," *Speech Communication*, vol. 50, no. 1, pp. 38–55, 2008.
- [18] Okko Räsänen, Unto K Laine, and Toomas Altsaar, "Blind segmentation of speech using non-linear filtering methods," *Speech Technologies*, pp. 105–124, 2011.
- [19] Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer, and Dan Chazan, "Phoneme alignment based on discriminative learning," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [20] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldi.," in *Interspeech*, 2017, pp. 498–502.
- [21] Sarah King and Mark Hasegawa-Johnson, "Accurate speech segmentation by mimicking human auditory processing," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 8096–8100.

- [22] Joerg Franke, Markus Mueller, Fatima Hamlaoui, Sebastian Stueker, and Alex Waibel, "Phoneme boundary detection using deep bidirectional lstms," in *Speech Communication; 12. ITG Symposium*. VDE, 2016, pp. 1–5.
- [23] Liang Lu, Lingpeng Kong, Chris Dyer, Noah A Smith, and Steve Renals, "Segmental recurrent neural networks for end-to-end speech recognition," *arXiv preprint arXiv:1603.00223*, 2016.
- [24] Yossi Adi and Joseph Keshet, "Structured: risk minimization in structured prediction," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2282–2286, 2016.
- [25] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang, "A tutorial on energy-based learning," *Predicting structured data*, vol. 1, no. 0, 2006.
- [26] O. J. Räsänen, U. K. Laine, and T. Altsaar, "An improved speech segmentation quality measure: the r-value," in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [27] A. Ben-Shalom, J. Keshet, D. Modan, and A. Laufer, "Automatic tools for analyzing spoken hebrew," in *Afeka Conference for Speech Processing*, 2014.