This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Seshadri, Shreyas; Räsänen, Okko

SylNet

Published in:
IEEE Signal Processing Letters

DOI:
[10.1109/LSP.2019.2929415](https://doi.org/10.1109/LSP.2019.2929415)

Published: 01/09/2019

Document Version
Peer reviewed version

Please cite the original version:
Seshadri, S., & Räsänen, O. (2019). SylNet: An Adaptable End-to-End Syllable Count Estimator for Speech. *IEEE Signal Processing Letters*, 26(9), 1359-1363. <https://doi.org/10.1109/LSP.2019.2929415>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

SylNet: An Adaptable End-to-End Syllable Count Estimator for Speech

Shreyas Seshadri and Okko Räsänen

Abstract—Automatic syllable count estimation (SCE) is used in a variety of applications ranging from speaking rate estimation to detecting social activity from wearable microphones or developmental research concerned with quantifying speech heard by language-learning children in different environments. The majority of previously utilized SCE methods have relied on heuristic DSP methods, and only a small number of bi-directional long short-term memory (BLSTM) approaches have made use of modern machine learning approaches in the SCE task. This paper presents a novel end-to-end method called SylNet for automatic syllable counting from speech, built on the basis of a recent developments in neural network architectures. We describe how the entire model can be optimized directly to minimize SCE error on the training data without annotations aligned at the syllable level, and how it can be adapted to new languages using limited speech data with known syllable counts. Experiments on several different languages reveal that SylNet generalizes to languages beyond its training data and further improves with adaptation. It also outperforms several previously proposed methods for syllabification, including end-to-end BLSTMs.

Index Terms—syllable count estimation, end-to-end learning, deep learning, speech processing.

I. INTRODUCTION

Automatic syllable count estimation (SCE) from speech is used in a number of applications. For instance, syllable-based speaking rate estimation algorithms (such as [1], [2]) can be used to analyze prosodic patterns of speakers and speaking styles for linguistic research, or used as additional information for training text-to-speech (TTS) synthesis systems. Syllables are also used for automatic estimation of vocal activity and social interaction from long and noisy audio recordings captured by wearable microphones, as in the personal life log application of Ziaei et al. [3] [4]. There is also a need for robust language-independent methods for quantifying the amount of speech in daylong child-centered audio recordings from various language environments [5], [6], as child language researchers use such data to understand language development in children in (e.g., [7], [8]).

The main motivation for choosing SCE over automatic speech recognition (ASR) in these tasks comes from the robustness and potential language-independence of the syllabification algorithms. While ASR may be applicable in high-resource domains and good recording conditions, deployment of reliable ASR systems across a large set of (potentially low-resource) languages is still challenging in a manner that also

ensures consistent performance in the adverse signal conditions typical to daylong recordings from wearable devices. In contrast, syllables are fundamentally based on the rhythmic succession of low- and high sonority speech sounds present in all languages. Since sonority is closely related to the amplitude modulation of speech signals, with syllabic nuclei being the locally most energetic speech sounds (see, e.g., [9] or [10] for an overview), syllables are potentially robust against noisy environments and channel variations.

However, so-far the most widely used syllabifiers (segment boundary and nuclei detectors) have been based on heuristic DSP operations [1] [11] [2] [10] [12] [13] [14]. According to our knowledge, only the bi-directional long short-term memory (BLSTM) -based algorithms by [6], [15], and [16] have used modern machine learning techniques to solve the SCE problem. In addition, the actual cross-language generalization capabilities of the algorithms, including [15] and [16], have largely remained untested, nearly all studies using speech data from only one language for development and testing.

This paper addresses the gaps in the literature by proposing a new method called SylNet for automatic SCE from speech, making use of state-of-the-art WaveNet-like [17] architecture for end-to-end SCE. We benchmark SylNet and a number of existing unsupervised and supervised algorithms on three languages different from the training data. We also investigate adaptability of different approaches for improved performance when a small amount of speech with known syllable counts is available in the target language. Unlike the previous BLSTM-models by [15] and [6] and similar to [16], the newly proposed SylNet does not require temporally aligned syllable- and phone-level annotations for training or adaptation. As a result, we show that SylNet achieves state-of-the-art SCE performance on novel languages without any language-specific tuning, and can be improved even further with adaptation.

II. SYLNET

Figure 1 shows the structure of the SylNet model¹, which is derived from the recently developed WaveNet [17] model used for speech synthesis. Instead of autoregressive waveform sample prediction as in speech synthesis, we apply the model, along with an additional LSTM layer, to the prediction of accumulated syllable counts based on spectral features of speech. The motivation for using a WaveNet-like architecture lies in the efficient training and use of multi-scale information in speech signals and easy adaptability of the network

Shreyas Seshadri is with the Department of Signal Processing and Acoustics, Aalto University, Finland. e-mail: (shreyas.seshadri@aalto.fi)

Okko Räsänen is with the Faculty of Information Technology and Communication Sciences, Tampere University, Finland, and the Department of Signal Processing and Acoustics, Aalto University, Finland.

¹Python + TensorFlow implementation of SylNet with a pre-trained language-independent model is available for download at <https://github.com/shreyas253/SylNet>.

Fig. 1. Structure of the SylNet model with input layer on the left and the rest of the N layers on the right. Each of the N layers consists of K channels, all feeding a shared PostNet layer that integrates information across the entire input using an LSTM, and then to a dense layer that maps activations from all scales and channels into one scalar number (when using L1 loss in Eq. 1) or vector (ordinal loss in Eq. 2) corresponding to estimated syllable count.

through tuning only the parameters responsible for merging information from different feature extraction layers.

Input to the model consists of time-frequency features of speech sampled at a regular interval (T timesteps, D dimensions). The input layer has K -channels each consisting of a w -point gated convolutional unit. This is followed by N layers, each consisting of similar gated convolutional units as the input layer. The outputs of these gated convolutional units are passed through an affine transform and added with a residual [18] of the original input, and then passed to the next layer. Skip connections [19] pass the outputs of the gated convolutional unit from each of the N layers to a so-called PostNet layer through separate layer-specific affine transformations. The residual [18] and skip [19] connections help in preventing the problem of diminishing gradients in deeper networks by allowing the error gradients to pass through, thus reducing convergence time and enabling the use of deep models [17]. The PostNet layer integrates the information coming from the skip connections of all the N layers through addition. This sum is then fed to a rectified linear unit (ReLU) after a w -point affine convolution, to an LSTM layer that accumulates syllable count information over time, and finally to a dense layer with 1 or R units (depending on the loss; Eqs. 1 and 2).

During training, only the PostNet outputs of the last time frame for each utterance are used in the calculation of the loss function, which compares estimated syllable counts to the known ground truth for that utterance (Eqs. 1 and 2). As in any deep convolutional architecture, the N hidden layers can be considered as feature extractors that compute time-frequency patterns of an increasing non-linearity and complexity that are relevant for the current task. The PostNet layer does the job of assimilating these features to the syllable count. This structure has the advantage that a pre-trained model can be adapted to another (potentially much smaller) dataset in different language by only re-training the PostNet layers that only have a limited number of parameters. As we show in the experiments, this avoids overfitting even when very little adaptation data are available.

III. COMPARED SCE METHODS

Our experiments in Section V compare the SylNet to a number of alternatives from two classes of SCE algorithms: so-called envelope-based methods and end-to-end methods, where the latter category also includes the present SylNet approach.

A. Envelope prediction approaches

All envelope-based SCE methods consist of two basic stages (see Fig. 2): 1) computation of a so-called sonority envelope for the speech input, in which temporally local peaks are assumed to correspond to syllabic nuclei in the input, and 2) automatic peak picking to extract nuclei count from the envelope. Optionally, a third stage consisting of a least-squares linear model from model outputs to reference syllable counts (or to speaking rate as in [2], or word counts, as in [4], [6], [20]) can be used, allowing correction of any systematic biases in the estimation procedure. The envelope estimation can be either based on a series of DSP operations (hence called a “blind” or “unsupervised” approach), or by a supervised machine-learning approach to map incoming speech into syllabic nuclei likelihoods.

In this study, two previously reported unsupervised methods were compared: “**thetaSeg**” that carries out perceptually motivated entrainment to sonority fluctuations in speech with proven cross-language capability for syllable segmentation [10], and the envelope detector stage from speaking rate estimator by Wang and Narayanan (“**WN**”, [2]) that was also found to be the best syllabifier for word count estimation in [4]. In addition, we include a supervised BLSTM-based system described in [20], developed further from the one reported in [15], here referred to as **BLSTM-env**. The first two consist of a series of DSP operations to transform signal waveform $\mathbf{x}_u$ of an utterance u into a smooth sonority envelope $\mathbf{y}_u$ to be used for peak picking. In contrast, BLSTM-env consists of two stacked BLSTM layers, each having 60 LSTM cells, that are responsible for converting 24-channel log-Mel (25-ms window with 10-ms time steps) input features into syllable nucleus likelihoods $\mathbf{y}_u$ as a function of time, again followed by peak picking. In BLSTM-env training, nucleus targets are presented as Gaussian-shaped kernels centered around the nucleus phone of each manually annotated syllable (see [15]). The basic structure of these methods is elaborated in Figure 2, and more details can be found in the respective papers ([10] [2], [20]).

For all three envelope-based approaches, the initial syllable counts n_u are acquired with a peak picking algorithm that uses a threshold hyperparameter θ , which determines the minimum required amplitude difference between a local maximum and the previous local minimum in order for the maximum to be considered as a nucleus. These counts are then mapped to final syllable count estimates using a linear mapping $s_u = \alpha n_u + \beta$. During training θ , α , and β are jointly estimated using an exhaustive search over all possible thresholds θ with their corresponding least-squares linear models, choosing the triplet that the minimizes relative error in Eq. (1) on the training data. We did not optimize the large number of internal parameters (filter lengths etc.) of the unsupervised methods, but used the ones found to be optimal in the original papers [2] [10].

B. End-to-end methods

End-to-end methods aim to directly predict the syllable count from speech utterances without any hand-specified intermediate representational stages. We examine two supervised

Fig. 2. Block diagrams depicting the basic structure of the envelope estimation methods used for syllable count estimation (adapted from [20]).

neural network methods for end-to-end SCE that both use 24-channel log-Mel speech features as inputs (as in the BLSTM-env; see Fig. 3): 1) **BLSTM-count**, which is a BLSTM network having five hidden LSTM layers (2 stacked standard BLSTMs + a forward LSTM layer with linear activation), i.e., a similar structure to the BLSTM-env, but now having an additional forward LSTM layer to output estimated syllable counts at the end of the input utterance (see also [16]). Each layer has 60 LSTM cells (per direction).

2) The newly proposed **SylNet**. After initial experimentation on training data (see Section IV), N and K were set to 10 and 128, respectively. Convolution kernel length w was set to 5, with the effective receptive field of the topmost layer corresponding to 53 frames. The final LSTM layer in the PostNet was also set to have 128 cells. The basic structure of the model was described in Section II, and is elaborated in Figures 1 and 3.

Fig. 3. Block diagrams depicting the basic structures of the end-to-end SCE methods for syllable count estimation.

Training Loss - We explored two different loss functions and associated output layers for SylNet training. The first loss term (also used by the BLSTM-count) was defined as the relative error between true and estimated syllable counts, as this metric was also used as the evaluation metric in the experiments due to its previous use in the evaluation of word-count estimation systems [4], [6], [20]. L1-based relative error was defined as:

$$e_{L1} = \frac{1}{M} \sum_{u=1}^M \frac{|\hat{s}_u - s_u|}{s_u} \quad (1)$$

where $\hat{s}_u$ and s_u correspond to the predicted and true target number of syllables for the M training utterances in each mini-batch respectively. With this loss, last layer of the network was single linear unit performing regression to $\hat{s}_u$.

However, the true syllable counts are discrete and ordinal in nature. Another potential loss function, earlier found to be more suited to ordinal data [21]–[23], treats the problem as a series of $R - 1$ binary classifications with R being the rank of the ordinal targets (set to the maximum number of syllables per utterance in the training data). Here the last layer of the DNN is defined as a set of $R - 1$ sigmoid units. The r th output represents the probability that the value of the output is larger than r . For example, target count of three would be represented as $[1 \ 1 \ 1 \ 0 \ \dots \ 0]$ with $R - 1 - 3$ zeros. Based on initial experiments, the relative ordinal objective loss between the sigmoid activations $\hat{o}_u$ and the true sigmoid targets o_u corresponding to the given reference count s_u was defined as:

$$e_{ord} = \frac{1}{M} \sum_{u=1}^M \frac{\|\hat{o}_u - o_u\|_2}{s_u} \quad (2)$$

and this was compared to the L1 loss in the experiments. For testing, $\hat{s}_u$ is obtained by thresholding sigmoid outputs at 0.5.

IV. EXPERIMENTAL SETUP

Our experimental setup aimed to test two conditions: 1) out-of-the-box generalization of the compared SCE algorithms to novel languages beyond the training data, and 2) performance on these novel languages when a limited amount of speech data with syllable counts (but not necessarily syllable positions) is available for algorithm adaptation.

A. Data

Two large speech corpora with syllable- and phone-level annotations were used in the for the initial training of the supervised syllable count estimation models and for optimization of detection thresholds and linear mapping coefficients of unsupervised methods: 1) **Phonetic Corpus of Estonian Spontaneous Speech** (EstPhon; [24]) and 2) **Korean Corpus of Spontaneous Speech** (Korean; [25]). The former includes several spontaneous Estonian dialogues with pairs of male and female talkers, totaling up to 10,158 utterances in high studio quality (5.2 hours of audio). Korean corpus also consists of conversational dialogues between talkers, including a wide range of speaker ages. In order to keep the Korean corpus comparable to EstPhon in size, a randomly sampled subset of 12,000 utterances (5.0 hours) was used for syllabifier training.

Three additional corpora with manually annotated syllable counts were chosen for testing and adaptation: 1) **FinDialogue Corpus of Spontaneous Finnish Speech** [26], having dialogues from 4 talkers (2 male; 64 minutes of speech in total), 2) **C-PROM corpus of spoken French** [27] consisting of 24 multi-minute recordings of various regional varieties of French from several discourse genres, totaling 1.2 hours of data, and 3) **Switchboard corpus of spontaneous American English telephone conversations** (SWB; [28]), using a syllable-annotated ICSI transcription project subset (hundreds of talkers, 159 minutes). While FinDiag is of high studio quality, C-PROM contains more varied recording conditions (e.g., including some signal clipping), whereas SWB consists of narrowband telephone landline recordings, hence causing a substantial channel mismatch between training and testing.

B. Training, Adaptation and Evaluation

SylNet and BLSTMs (with approximately 3M and 100k parameters, respectively) were trained on Korean and EstPhon with the ADAM optimizer [29] with hyperparameters β_1 and β_2 set to 0.9 and 0.999 and learning rate set to 10^{-4} , with a 50% dropout applied after every non-linearity, and an early stopping of 10 epochs. Minibatch size was set to 32 utterances. The test corpora were divided into separate adaptation and testing sets. Each test set was defined by randomly sampling approx. 50% of all utterances of the corpus for the set. To simulate varying amounts of adaptation data, 8 different-sized sets ranging from mere 30 seconds up to 45 minutes of speech were created from the remaining utterances, ensuring that no same speaker occurred in both adaptation and testing. Five randomly sampled sets of each of size were created, and average performance across these are reported in the results. For the envelope estimation methods, adaptation involved estimating the optimal peak picking parameter θ_L and the linear model parameters a_L and b_L that minimize relative syllable count error on the full adaptation set L . BLSTM-count was adapted by retraining the last LSTM layer, and SylNet by retraining the entire PostNet module. Similarly to [4], [6], [20], the methods were evaluated on the held-out test data using the L1-relative error (%) as shown in Eq. 1.

V. RESULTS

Figure 4 shows the results in terms of L1 relative error. On average, the end-to-end methods outperform the envelope-based approaches, even though BLSTM-count fails on FinDiag before adaptation. All methods generally improve performance as the size of the adaptation dataset increases. Overall, the SylNet produces the lowest relative error before adaptation and after full adaptation, even though WN and BLSTM-env have comparable or sometimes better performance on some C-PROM adaptation sets. Also, the effects of end-to-end system adaptation on FinDiag dataset are minimal (approx. $\sim 2\%$ points for SylNet). This is likely because Finnish has relatively similar phonetic and syllabic structure to Estonian used for model training. As for the SylNet loss functions, the ordinal loss (Eq. 2) outperforms the L1 loss both out-of-the-box and after adaptation.

In order to visualize how the SylNet accumulates syllable counts over time, Fig. 5 shows the PostNet output activations for an utterance taken from the FinDiag corpus together with the accumulating reference syllable counts or both the L1 and the ordinal losses. Since receptive field size of the deepest convolutional kernels is very large in time, the model utilizes information from both past and future frames in order to optimize the count estimate at the final frame. In both cases for the utterance shown, the final syllable count estimate is very accurate.

VI. CONCLUSIONS

This paper presented a novel end-to-end method called SylNet for syllable count estimation (SCE) from conversational speech and evaluated its performance on a number of languages, both out-of-the-box and by adapting the model to the

Fig. 4. L1-Relative errors on the SWB, FinDiag and C-PROM datasets for difference syllable count estimation methods with varied amount of adaptation data. Error bars correspond to one standard deviation across the 5 randomly sampled folds for adaptation data of each amount (see Section IV-B).

Fig. 5. An example of SylNet PostNet output accumulation (using L1 and ordinal loss) as a function of time on a Finnish utterance from FinDiag.

target language. The results show that the proposed algorithm generalizes well to novel languages, and further improves with adaptation data if syllable counts (but not necessarily temporal alignments) are available for some minutes of speech in the target language. In total, the work shows that supervised end-to-end systems can be utilized for language-independent SCE from speech, clearly outperforming the earlier DSP-based methods. Future work should seek ways to further improve robustness and generalization capabilities of SylNet, e.g., through various data augmentation techniques.

ACKNOWLEDGMENT

This study was funded by Academy of Finland grants no. 312105 and 314602. The authors thank Lauri Juvela for his help with the SylNet implementation.

REFERENCES

- [1] N. Morgan and E. Fosler-Lussier, "Combining multiple estimators of speaking rate," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-1998)*. IEEE, 1998, pp. 729–732.
- [2] D. Wang and S. S. Narayanan, "Robust speech rate estimation for spontaneous speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2190–2201, 2007.
- [3] A. Ziaei, A. Sangwan, L. Kaushik, and J. Hansen, "Prof-life-log: Analysis and classification of activities in daily audio streams," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-2015)*. IEEE, 2015, pp. 4719–4723.
- [4] A. Ziaei, A. Sangwan, and J. Hansen, "Effective word count estimation for long duration daily naturalistic audio recordings," *Speech Communication*, vol. 84, pp. 15–23, 2016.
- [5] "Analyzing child language experiences around the world -project," <https://sites.google.com/view/acledwid/home>, accessed: 2019-01-28.
- [6] O. Räsänen, S. Seshadri, J. Karadayi, E. Riebling, J. Bunce, A. Cristia, F. Metze, M. Casillas, C. Rosemberg, E. Bergelson, and M. Soderstrom, "Automatic word count estimation from daylong child-centered recordings in various language environments using language-independent syllabification of speech," *submitted for publication*.
- [7] E. Bergelson, A. Amatuni, S. Dailey, S. Koorathota, and S. Tor, "What do north american babies hear? a large-scale cross-corpus analysis," *Developmental Science*, vol. e12715, 2018.
- [8] A. Weisleder and A. Fernald, "Talking to children matters: Early language experience strengthens processing and builds vocabulary," *Psychological Science*, vol. 24, pp. 2143–2152, 2013.
- [9] S. G. Parker, *Quantifying the sonority hierarchy*. PhD dissertation. Graduate School of the University of Massachusetts Amherst, MA., 2002.
- [10] O. Räsänen, G. Doyle, and M. C. Frank, "Pre-linguistic segmentation of speech into syllable-like units," *Cognition*, vol. 171, pp. 130–150, 2018.
- [11] N. Obin, F. Lamare, and A. Roebel, "Syll-o-matic: an adaptive time-frequency representation for the automatic segmentation of speech into syllables," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-2013)*. IEEE, 2013.
- [12] P. Mermelstein, "Automatic segmentation of speech into syllabic units," *Journal of the Acoustical Society of America*, vol. 58, pp. 880–883, 1975.
- [13] R. Villing, J. Timoney, T. Ward, and J. Costello, "Automatic blind syllable segmentation for continuous speech," in *Irish Signals and Systems Conference (ISSC-2004)*, 2004.
- [14] A. A. Reddy, N. Chennupati, and B. Yegnanarayana, "Syllable nuclei detection using perceptually significant features," in *Proc. Interspeech-2013*. ISCA, 2013, pp. 963–967.
- [15] C. Landsiedel, J. Edlund, F. Eyben, D. Neiberg, and B. Schuller, "Syllabification of conversational speech using bidirectional long-short-term memory neural networks," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-2011)*. IEEE, 2011, pp. 5256–5259.
- [16] Y. Jiao, V. Berisha, and J. Liss, "Online speaking rate estimation using recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-2016)*. IEEE, 2016, pp. 5245–5249.
- [17] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *9th ISCA Speech Synthesis Workshop*, 2016, pp. 125–125.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016, pp. 770–778.
- [19] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [20] O. Räsänen, S. Seshadri, and M. Casillas, "Comparison of syllabification algorithms and training strategies for robust word count estimation across different languages and recording conditions," in *Proc. Interspeech-2018*, 2018, pp. 1200–1204.
- [21] J. Cheng, Z. Wang, and G. Pollastri, "A neural network approach to ordinal regression," in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. IEEE, 2008, pp. 1279–1284.
- [22] Z. Niu, M. Zhou, L. Wang, X. Gao, and G. Hua, "Ordinal regression with multiple output CNN for age estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4920–4928.
- [23] P. A. Gutierrez, M. Perez-Ortiz, J. Sanchez-Monedero, F. Fernandez-Navarro, and C. Hervás-Martínez, "Ordinal regression methods: survey and experimental study," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 1, pp. 127–146, 2016.
- [24] P. Lippus, T. Tuisk, N. Salveste, and P. Teras, "Phonetic corpus of estonian spontaneous speech," *Institute of Estonian and General Linguistics, University of Tartu*. DOI: <https://doi.org/10.15155/TY.D>, 2013.
- [25] W. Yun, K. Yoon, S. Park, J. Lee, S. Cho, D. Kang, K. Byun, H. Hahn, and J. Kim, "The korean corpus of spontaneous speech," *Daegu: Industry-Academic Cooperation Foundation, Keimyung University (Distributor)*, 2015.
- [26] M. Lennes, "Segmental features in spontaneous and read-aloud finnish," *Phonetics of Russian and Finnish. General Introduction. Spontaneous and Read-Aloud Speech*, V. de Silva and R. Ullakonoja, Eds. Peter Lang GmbH, pp. 145–166, 2009.
- [27] M. Avanzi, A.-C. Simon, J.-P. Goldman, and A. Auchlin, "C-prom: An annotated corpus for french prominence study," in *Speech Prosody 2010-Fifth International Conference*, 2010.
- [28] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "Switchboard: Telephone speech corpus for research and development," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-1992)*, vol. 1. IEEE, 1992, pp. 517–520.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.