



HAL
open science

GCI DETECTION FROM RAW SPEECH USING A FULLY-CONVOLUTIONAL NETWORK

Luc Ardaillon, Axel Roebel

► **To cite this version:**

Luc Ardaillon, Axel Roebel. GCI DETECTION FROM RAW SPEECH USING A FULLY-CONVOLUTIONAL NETWORK. 2019. hal-02438881

HAL Id: hal-02438881

<https://hal.science/hal-02438881>

Preprint submitted on 14 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GCI DETECTION FROM RAW SPEECH USING A FULLY-CONVOLUTIONAL NETWORK

Luc Ardaillon, Axel Roebel

UMR 9912 STMS (IRCAM / Sorbonne University / CNRS), Paris, France

ABSTRACT

Glottal Closure Instants (GCI) detection consists in automatically detecting temporal locations of most significant excitation of the vocal tract from the speech signal. It is used in many speech analysis and processing applications, and various algorithms have been proposed for this purpose. Recently, new approaches using convolutional neural networks have emerged, with encouraging results. Following this trend, we propose a simple approach that performs a regression from the speech waveform to a target signal from which the GCI are easily obtained by peak-picking. However, the ground truth GCI used for training and evaluation are usually extracted from EGG signals, which are not reliable and often not available. To overcome this problem, we propose to train our network on high-quality synthetic speech with perfect ground truth. The performances of the proposed algorithm are compared with three other state-of-the-art approaches using publicly available datasets, and the impact of using controlled synthetic or real speech signals in the training stage is investigated. The experimental results demonstrate that the proposed method obtains similar or better results than other state-of-the-art algorithms and that using large synthetic datasets with many speaker offers better generalization ability than using a smaller database of real speech and EGG signals.

Index Terms— GCI detection, epoch extraction, speech analysis, CNN

1. INTRODUCTION

Glottal Closure Instants (GCI) detection consists in finding the temporal locations of significant excitation of the vocal tract, that occur in voiced speech during the closure of vocal folds, directly from the speech signal. GCI detection has many applications in speech analysis and processing [1]. Examples of such applications are the precise estimation of the fundamental frequency (f_0) [2], analysis of vocal disorders [3, 4], formants estimation [5], or speech synthesis [6]. State-of-the-art f_0 estimators like [7] work very well in most cases, but are not well adapted to deal with rough voices that contain jitter or sub-harmonics for which the pitch is not well-defined. In such cases, using GCI detection instead of f_0 might improve the quality of resynthesized signals when used as a parameter in vocoders like SVLN [8] and PaN [9, 10, Section 3.5.2] for speech analysis/resynthesis and transformation. In this direction, we proposed in [10, Section 6.3.3] to extract pitch marks from rough singing voice extracts to estimate jitter and shimmer and use it to synthesise singing voices with rough characteristics. A good GCI detection would thus help to better synthesize and transform such types of voices that are currently not well handled.

One possible way to extract the GCI positions in a speech signal is to use Electroglottographic (EGG) recordings. Electroglottography uses dedicated hardware to record the vocal folds contact area by placing two contact electrodes on the speakers neck. GCIs can

then be extracted using peak-picking on the derivative of the EGG signal. However, such recordings are rarely available, which raised the need for methods that can extract GCIs directly from speech signals. Many algorithms have been proposed for this purpose. Until recently, all approaches used to be based on hand-crafted digital signal processing techniques and heuristics. Thorough reviews of those techniques can be read in [11, 12, 1], where authors compared their performances on a variety of speech and singing signals. Typically, such methods first compute an intermediate speech representation, such as the linear prediction residual [13], a zero-frequency filtered signal [14], or a mean-based signal [15], which emphasizes the locations of glottal closure instants found at local maxima, impulses or discontinuities. Then, dynamic-programming or peak-picking is used to select the GCIs among the detected candidates. Examples of such approaches are the SEDREAMS [15], DYPSA [13], DPI [16], YAGA [17], ZFR [14], or MMF [18] algorithms. Although such approaches have been shown to perform reasonably well, they rely on different processing techniques that require manual tuning of parameters (e.g. the mean f_0 value for SEDREAMS [15]), and the quality of their results remains quite dependant on the characteristics of the analyzed speech signal (e.g. pitch and voice quality, speech or singing voice) and the targeted application [11, 12]. In particular, it has been noted in [12] that all tested algorithms tend to perform better on speech than on singing. Moreover, some algorithms like SEDREAMS [15] or DYPSA [13] also detect some GCIs during unvoiced segments and thus rely on further algorithms to filter out GCI candidates in unvoiced parts.

To overcome the limits of previous methods, new data-driven approaches have been recently proposed. In [19], authors used extremely randomized trees (ERT) to classify peaks from the speech waveforms as being a GCI or not. Other recent studies proposed to use convolutional neural networks [20, 3, 21]. Similarly to [19], the authors in [20] used a classification approach where GCI candidates are the negative peaks of a low-pass filtered signal. Similarly, [3] also employed a classification-based approach using 3 parallel CNNs operating on different signal representations (including the linear prediction residual). In [21], the authors used a CNN to optimize both a classification and regression cost, where a GCI is simultaneously detected and localized in a frame. Other recent related works used regression-based approaches with neural networks for f_0 [22] or glottal source parameters estimation (including GCI) [23]. However, those approaches all rely on EGG signals for establishing the ground truth reference used for training the networks. This has two main drawbacks: 1) the EGG signals are often noisy, and the extracted ground-truth GCIs are thus likely to contain errors; 2) EGG signals are rarely available, which makes it difficult to build large multi-speaker databases for training, and thus may limit the generalization ability of the models.

Building up on previous work on pitch estimation of speech signals [7] we propose here to use a simple and efficient Fully-Convolutional Network (FCN) to perform a mapping between the

speech waveform and a target signal from which the GCIs can be easily extracted using peak-picking. Similarly to other methods, our model can be trained on real signals when EGG recordings are available. But in order to avoid the drawbacks of relying on EGG signals, we also propose to train our network on a database of high-quality synthetic speech signals with a perfectly reliable ground truth. The performances of the proposed algorithm have been compared with three other state-of-the-art approaches using publicly available datasets, and the impact of using controlled synthetic speech or real speech signals in the training stage is explored.

In section 2, we will present an overview of the proposed approach and detail the architecture of our network. Then we will present in section 3 the datasets that have been used in our experiments. Finally, the methodology and results of our evaluations will be presented in section 4.

2. PROPOSED APPROACH

2.1. overview

In this paper, we propose to use a fully-convolutional network (FCN) that takes a raw speech waveform as input to perform a regression to a simple 1-dimensional target waveform from which the GCI positions can be easily extracted by means of pick-peaking. The network is trained to minimize the mean-square error between the target and predicted values. The following sections detail the target signals used for the regression, the network architecture, and the procedure for GCI extraction. A python implementation of our approach, along with pretrained models, is made available online¹.

2.2. Target signals

An obvious target signal for our purpose would be a train of diracs placed at GCI positions. Initial experiments showed, however, that it is difficult for a neural network to learn such a representation, since a small variation in the input waveform (i.e. a shift by one sample) leads to a drastic change in the target. We thus need to use target signals with smoother variations. We propose to use as a target shape a triangle whose maximum of value 1 is placed at the GCI position and whose miminas of value 0 are placed at half a period to the left and right of the GCI (based on an f_0 analysis). However, besides the GCI position, this shape is not correlated to other source parameters. Using synthetic signals, a particular advantage is that the source parameters are controlled (in our case the LF model parameters [24]. See section 3.2). Thus, any other shape relying on those parameters might be used, the most obvious choice being the glottal flow itself. We thus propose, as an alternative to triangles, to use the glottal flow as a target shape when training our model with synthetic signals. In order to be independent of the signal's energy, each period is normalized with a maximum of 1. Figure 1 illustrates an example of a synthetic speech waveform with the corresponding ground truth GCIs and both the triangle and glottal flow target signals.

2.3. Network architecture

The proposed network architecture, composed of 7 convolutional layers, is similar to the one used in [7]. Max pooling is included after the first 3 convolutional layers, each layer is followed by batch normalization, and their output are passed through relu activations, except for the last layer for which sigmoid activations are used. The number of filters in each layer and their sizes are detailed in figure

2. All convolutions are valid (no padding applied) with a stride of 1. The minimal input to obtain 1 output value is a 993-samples excerpt of a speech signal sampled at 16 kHz, but any sound of size > 993 might be used as input. Note that the same size was used in [7], but with a sampling rate of 8kHz. Then, the only differences with [7] are the number of filters in the convolutional layers 1 to 6 that have been doubled, and the last convolutional layer that contains only 1 filter instead of 486.

2.4. GCI extraction

Since the network is trained on 16kHz speech signals, the predicted target at the output of the network has a sampling rate of only 2kHz, because of the down-sampling performed by the 3 max-pooling layers. For better precision, the predicted target is thus up-sampled back to 16kHz using cubic spline interpolation. Then, the GCIs are extracted using a peak-picking procedure that depends on the target signal. For the triangle target, a simple peak-picking with a minimal threshold of 0.5 is used. For the glottal flow target, the pick-peaking is performed on the negative peaks of its derivative, which coincide with the GCIs.

3. DATASETS

For the purpose of proper training and objective evaluations, we need a database with a reliable ground truth. As previously-exposed, publicly available databases with parallel speech and EGG recordings, such as the CMU artic one [25], are typically used for this purpose [20, 11, 16]. However, EGG recordings are often subject to low and high-frequency noise and distortions that happen during the recording. Then, depending on the quality of the signal, the pre-processing used, and the tuning of peak-picking parameters, the extracted GCIs are likely to contain many errors that will impact both the training and evaluation of algorithms.

To avoid these problems, we propose here to use re-synthesized speech signals, produced by means of a high-quality analysis/re-synthesis method achieving near transparent quality, for which the GCI positions are known and controlled (as was done in [7] for f_0 estimation). Using such signals ensures to have a perfectly reliable ground truth exempt of any errors. Another advantage of using a synthetic database is that one can easily include many different speakers with a wide variety of signal characteristics (speech or singing, pitch, voice quality, ...), whereas it is difficult to build a database with EGG recordings for many speakers.

3.1. CMU artic dataset

Concerning the training on real speech with EGG-derived ground truth GCIs, we used the publicly-available CMU artic database [25] with the 3 speakers BDL, JMK and SLT, (similarly to [21]). In order to obtain the ground truth GCIs, we first pre-processed the EGG signals using 5th-order high and low-pass filters with respective cut-off frequencies of 30 and 500Hz. Then, the ground-truth GCI positions were obtained using peak-picking on the negative peaks of the derivative of the pre-processed EGG signal (as done in [12, 21]). The pick-peaking was done using the code provided by the authors of [21], available online² (getpeaks function in dataloader.py) to have a comparable ground truth for evaluation. Then, the extracted GCI positions were used to build the corresponding triangle target signals. Note that the glottal flow target cannot be used in this case due to the missing information about the glottal pulse parameters.

¹https://github.com/ardaillon/FCN_GCI

²<https://github.com/VarunSrivastavaIITD/DCNN>

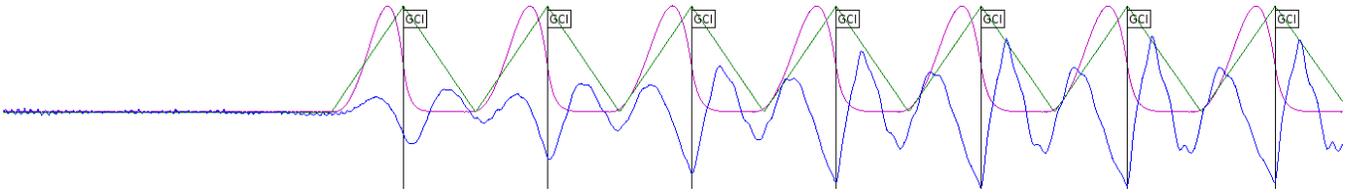


Fig. 1. Example of a synthesized speech waveform (blue) with the corresponding GCI markers (black) and corresponding triangle (green) and glottal flow (purple) target signals

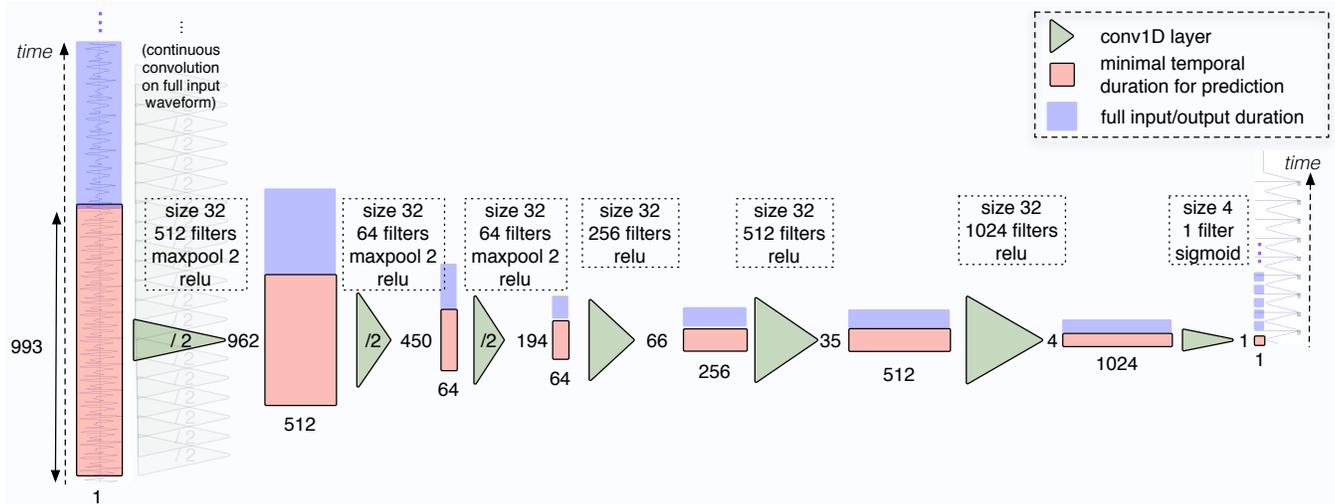


Fig. 2. Architecture of the proposed fully-convolutional network. All convolutions are valid with stride 1. Minimum input size is 993.

3.2. Synthetic dataset

For our synthetic dataset, we used the recent PaN vocoder described in [10, Section 3.5.2][9] to resynthesize existing speech signals in a controlled way with a near-transparent quality. Basically, the PaN vocoder uses the f_0 and R_d parameter contours analysed on the original speech signal in order to generate a sequence of pulses based on the LF model of glottal source [24], in an analysis/re-synthesis process. Then, the vocal tract filter is applied on the pulses and the original unvoiced component is added to the voiced source to generate the complete signal. In order to create this dataset, we first merged the publicly-available BREF [26] and TIMIT [27] datasets, ending up into a total of 11616 short sentences spoken by respectively 100 French and about 630 English speakers. Then, each sentence has been re-synthesized using the PaN synthesis engine. In order to augment the diversity of source shapes seen by the network during the training, and thus better represent the diversity that may be encountered in real speech recordings, we also re-synthesized the same signals with R_d contours shifted by fixed values of $+0.5$ and -0.5 . During the synthesis process, we also store the GCI positions used to produce the ground truth and the triangle target, as well as the normalized glottal flow target signals.

4. EVALUATION

4.1. Methodology

Two state of the art algorithms for GCI detection are SEDREAMS [11] and DPI [16, 21]. We thus chose to compare the results of our

approach to those 2 algorithms. Additionally, we also use the results given for the DCNN approach in [21] as another reference in our evaluation. To evaluate our approach in the same conditions than DCNN [21], a first version of our model, called "FCN-CMU-10/90", has been trained on the CMU artic database using a 10/90 train/test split, where the test split is also used for validation. However, since the resulting training set is very small, we also trained a 2^{nd} version with more training data, called "FCN-CMU-60/20/20", using a 60/20/20 training/validation/test split. As the glottal flow signal is not accessible for real signals, only the triangle target shape is used in this case. Then, 2 additional versions of our models, called "FCN-synth-tri" and "FCN-synth-GF" have been trained on our synthetic database using respectively the triangle and the glottal flow shapes as target signals, using a 60/20/20 split in both cases.

Then, all trained models, as well as the SEDREAMS and DPI algorithms are evaluated on the tests splits of both the synthetic and CMU artic datasets. Additionally, in order to assess the generalization ability of the trained model on unseen speakers, we also evaluate the algorithms on 2 speakers (M01 and F01) from the PTDB-TUG database [28].

For the SEDREAMS algorithm, we used the implementation available in the COVAREP repository³ [29]. The provided function requires the mean f_0 value as a parameter. For the CMU and PTDB databases, we thus used a value adapted to each speaker. However, for the synthetic database which contains many speakers, we used a single value of 150Hz. For the DPI algorithm, we used an implementation provided to us by the authors of [21].

³<http://covarep.github.io/covarep/>

Table 1. Results of evaluations on various databases using standard metrics (DCNN results have been copied from [21])

	IDR			MR			FAR			IDA		
	synth	CMU	PTDB	synth	CMU	PTDB	synth	CMU	PTDB	synth	CMU	PTDB
FCN-synth-tri	99.90	97.95	95.37	0.08	1.89	3.40	0.02	0.17	1.22	0.08	0.26	0.32
FCN-synth-GF	99.91	98.43	95.64	0.06	1.20	2.91	0.04	0.37	1.45	0.11	0.34	0.38
FCN-CMU-10/90	49.63	99.39	90.13	48.05	0.50	8.91	0.51	0.11	0.95	0.52	0.10	0.26
FCN-CMU-60/20/20	60.06	99.52	88.17	39.14	0.40	11.00	0.64	0.08	0.81	0.50	0.09	0.26
SEDREAMS	89.26	99.04	95.34	3.86	0.21	2.15	6.88	0.75	2.51	0.68	0.36	0.62
DPI	88.22	98.69	91.3	2.14	0.23	2.16	9.64	1.08	6.53	0.83	0.23	0.49
DCNN (from [21])		99.3			0.3			0.4			0.2	

Table 2. results on CMU with modified metrics considering all GCIs

	IDR	MR	FAR	IDA
FCN-synth-GF	92.73	2.86	10.87	0.47
SEDREAMS	90.74	0.19	61.38	0.30

4.2. Training procedure

In all of our experiments, the networks are trained to minimize a mean-square error loss using mini-batch gradient descent with the Adam optimizer [30], with mini-batches composed of 128 examples randomly selected from the training set. An initial learning rate of 0.0002 is used, and a reduction by a factor 0.75 is applied when the validation loss doesn't decrease for 10 epochs (where an epoch consists of 500 batches), with a minimum value 0.000025. The training is stopped when the validation accuracy has not improved for 64 epochs. The input segments used for training had a fixed size of 993 samples. Contrary to [7], no normalization is applied on the input segments, but a global normalization to a maximum level of -3dB is applied offline on each file before training and evaluation.

4.3. Results

The considered algorithms are evaluated both in terms of reliability and accuracy using standard metrics defined in [31]: Identification Rate (IDR - % of correct detections, higher is better), Miss Rate (MR - % of missed detections, lower is better), False Alarm Rate (FAR - % of false insertions, lower is better) and Identification Accuracy (IDA - standard deviation of distance between the true and predicted GCIs, lower is better). Table 1 summarizes the results of our evaluation for all algorithms and databases. As can be observed, at least one of our models obtains the top results in all cases, except for the MR metric. In particular, our model "FCN-CMU-10/90" obtains slightly better results than the values reported in [21] for the DCNN model for similar evaluation settings. As could be expected, the models obtain their best results on the database they have been trained on. The models trained on synthetic data ("FCN-synth") also work well on real signals, but the inverse is not true. A probable explanation for this is that the use of many speakers in the synthetic database allows to better generalize the results to unseen speakers than for the FCN-CMU models which were trained on only 3 speakers. This tends to be confirmed by the better IDR and MR results of the FCN-synth models on the PTDB database. Regarding the FCN-synth models, the use of a glottal flow target resulted in better IDR and MR values than for triangles. However, SEDREAMS and DPI still exhibit better IDR performances than our FCN-synth models on the CMU database. (Note that the better performances of the SEDREAMS algorithm here compared to the values reported in [21] might be explained by the use of a different implementation or a better tuning of the mean f_0 parameter).

But the metrics used in table 4 (based on code from [21] following definitions from [13]) only consider ground truth GCIs in voiced parts of speech for which the left and right periods correspond to a range of [50-500]Hz. This thus excludes isolated pulses as well as GCIs at the extremities of voiced segments, which correspond to real glottal activity but are harder to properly detect [20]. The used metrics don't account neither for the actual number of wrongly-detected GCIs that are not contained within one period around a ground-truth GCI. These limitations thus result in better evaluation values compared to considering the ground truth as a whole, and don't fully account for the real performances of the algorithms. As an example, we show in table 2 the results of our evaluation on the CMU artic database using an alternative implementation of the metrics that considers all of the ground truth GCIs, regardless of the f_0 and voicing. As can be observed, this greatly impacts the results and in those conditions, our model trained on synthetic data gets better IDR and FAR measures than the SEDREAMS algorithm. The very high FAR value for SEDREAMS in this case is due to the high number of wrongly-detected GCIs output by the algorithm during unvoiced segments, that were not considered in the results of table 4.

5. CONCLUSION

In this paper, we proposed a novel approach to GCI detection using a fully-convolutional network and a simple peak-picking procedure. In order to avoid the problems of using EGG signals, we proposed to train our model using high-quality synthetic speech. The results of our evaluations on publicly-available datasets showed that our approach performed similarly or better than other state-of-the-art algorithms in many cases, and that using a large database of synthetic signals with many speakers for training offers better generalization abilities than using a small database of real signals with EGG recordings of few speakers. In future works, more variations, like jitter or shimmer may be included into the synthetic dataset in order to better cover all possible variations encountered in real voice signals. Besides the GCI detection problem, the prediction of the glottal flow by our model might also allow to extract other useful voice source parameters.

6. ACKNOWLEDGEMENT

The authors would like to thank Dr. Goyal, Mohit, Dr. Srivastava, Varun and Dr. Prathosh, AP for their valuable help in providing matlab and python codes used in this article for establishing the ground truth GCI from the EGG recordings and for computing the metrics, as well as their implementation of the DPI algorithm.

7. REFERENCES

- [1] Thomas Drugman, Paavo Alku, Abeer Alwan, and Bayya Yegnanarayana, “Glottal source processing: From analysis to applications,” *Computer Speech and Language*, vol. 28, no. 5, pp. 1117–1138, 2014.
- [2] B. Yegnanarayana and K. Sri Rama Murty, “Event-based instantaneous fundamental frequency estimation from speech signals,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 17, no. 4, pp. 614–624, 2009.
- [3] Gurunath Reddy M, Tanumay Mandal, and Krothapalli Sreenivasa Rao, “Glottal Closure Instants Detection From Pathological Acoustic Speech Signal Using Deep Learning,” 2018.
- [4] Pranav S. Deshpande and M. Sabarimalai Manikandan, “Effective Glottal Instant Detection and Electroglottographic Parameter Extraction for Automated Voice Pathology Assessment,” *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 2, pp. 398–408, 2018.
- [5] Joseph M. Anand, S. Guruprasad, and B. Yegnanarayana, “Extracting formants from short segments of speech using group delay functions,” *INTERSPEECH*, vol. 2, no. 1, pp. 1009–1012, 2006.
- [6] Thomas Drugman and Thierry Dutoit, “The Deterministic plus Stochastic model of the residual signal and its applications,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 3, pp. 968–981, 2012.
- [7] Luc Ardaillon and Axel Roebel, “Fully-Convolutional Network for Pitch Estimation of Speech Signals,” pp. 0–4.
- [8] Gilles Degottex, Pierre Lanchantin, Axel Roebel, and Xavier Rodet, “Mixed source model and its adapted vocal tract filter estimate for voice transformation and synthesis,” *Speech Communication*, vol. 55, no. 2, pp. 278–294, 2013.
- [9] Stefan Huber and Axel Roebel, “On glottal source shape parameter transformation using a novel deterministic and stochastic speech analysis and synthesis system,” in *Interspeech 2015*, 2015, number 1, pp. 289–293.
- [10] Luc Ardaillon, *Synthesis and expressive transformation of singing voice*, Ph.D. thesis, EDITE; UPMC-Paris 6 Sorbonne Universités, 2017.
- [11] Thomas Drugman, Mark Thomas, Jon Gudnason, Patrick Naylor, and Thierry Dutoit, “Detection of glottal closure instants from speech signals: A quantitative review,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 3, pp. 994–1006, 2012.
- [12] Onur Babacan, Thomas Drugman, Nicolas D’Alessandro, Nathalie Henrich, and Thierry Dutoit, “A quantitative comparison of glottal closure instant estimation algorithms on a large variety of singing sounds,” *INTERSPEECH*, , no. August, pp. 1702–1706, 2013.
- [13] Patrick A Naylor, Anastasis Kounoudes, Jon Gudnason, and Mike Brookes, “Estimation of Glottal Closure Instants in Voiced Speech Using the DYPSA Algorithm,” vol. 15, no. 1, pp. 34–43, 2007.
- [14] K Sri Rama Murty and B Yegnanarayana, “Epoch Extraction From Speech Signals,” vol. 16, no. 8, pp. 1602–1613, 2008.
- [15] Thomas Drugman and Thierry Dutoit, “Glottal Closure and Opening Instant Detection from Speech Signals,” in *Interspeech 2009*, 2009.
- [16] A P Prathosh, T V Ananthapadmanabha, A G Ramakrishnan, and Senior Member, “Epoch Extraction Based on Integrated Linear Prediction Residual Using Plosion Index,” vol. 21, no. 12, pp. 2471–2480, 2013.
- [17] Mark R.P. Thomas, Patrick A. Naylor, and Jon Gudnason, “Estimation of Glottal Closing and Opening Instants in Voiced Speech Using the YAGA Algorithm,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 1, pp. 82–91, 2012.
- [18] Vahid Khanagha, Khalid Daoudi, and Hussein Yahia, “Detection of Glottal Closure Instants based on the Microcanonical Multiscale Formalism,” 2014.
- [19] Jindich Matousek and Daniel Tihelka, “Classification-based detection of glottal closure instants from speech signals,” in *INTERSPEECH*, 2017, vol. 2017-August, pp. 3053–3057.
- [20] Shuai Yang, Zhiyong Wu, Binbin Shen, and Helen Meng, “Detection of glottal closure instants from speech signals: A convolutional neural network based method,” in *INTERSPEECH*, 2018, vol. 2018-Sept, pp. 317–321.
- [21] Mohit Goyal, Varun Srivastava, and Prathosh AP, “Detection of Glottal Closure Instants from Raw Speech using Convolutional Neural Networks,” 2019.
- [22] Akihiro Kato and Tomi Kinnunen, “Waveform to Single Sinusoid Regression to Estimate the F0 Contour from Noisy Speech Using Recurrent Deep Neural Networks,” 2018.
- [23] A P Prathosh, Varun Srivastava, and Mayank Mishra, “Adversarial Approximate Inference for Speech to Electroglottograph Conversion,” pp. 1–14.
- [24] G. Fant, J. Liljencrants, and Q Lin, “Quarterly Progress and Status Report A four-parameter model of glottal flow,” *STL-Quarterly Progress and Status Report*, vol. 26, no. 4, pp. 1–13, 1985.
- [25] John Kominek and Alan W Black, “THE CMU ARCTIC SPEECH DATABASES,” in *5th ISCA Speech Synthesis Workshop*, 2004, number june 14-16, pp. 223–224.
- [26] J. L. Gauvain, Lori F Lamel, and M. Eskenazi, “Design Considerations and Text Selection for BREF, a large French Read-Speech Corpus,” *ICSLP*, , no. January 2013, pp. 1097–1100, 1990.
- [27] Victor Zue, Stephanie Seneff, and James Glass, “SPEECH DATABASE DEVELOPMENT AT MIT : TIMIT AND BEYOND,” vol. 9, pp. 351–356, 1990.
- [28] Gregor Pirker, Michael Wohlmayr, Stefan Petrik, and Franz Pernkopf, “A Pitch Tracking Corpus with Evaluation on Multipitch Tracking Scenario,” , no. August, pp. 1509–1512, 2011.
- [29] Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer, “COVAREP A COLLABORATIVE VOICE ANALYSIS REPOSITORY FOR SPEECH TECHNOLOGIES Computer Science Department , University of Crete , Heraklion , Greece Phonetics and Speech Laboratory , Trinity College Dublin , Ireland TCTS Lab - University of Mons , Belgium A,” *ICASSP*, pp. 960–964, 2014.
- [30] Diederik P. Kingma and Jimmy Ba, “Adam: A Method for Stochastic Optimization,” pp. 1–15, 2014.
- [31] Mark R.P. Thomas and Patrick A. Naylor, “The sigma algorithm: A glottal activity detector for electroglottographic signals,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 17, no. 8, pp. 1557–1566, 2009.