

# INTERPRETING GLOTTAL FLOW DYNAMICS FOR DETECTING COVID-19 FROM VOICE

Soham Deshmukh, Mahmoud Al Ismail, Rita Singh

Carnegie Mellon University, Pittsburgh, PA, USA  
 {sdeshmuk, mahmoudi, rsingh}@andrew.cmu.edu

## ABSTRACT

In the pathogenesis of COVID-19, impairment of respiratory functions is often one of the key symptoms. Studies show that in these cases, voice production is also adversely affected – vocal fold oscillations are asynchronous, asymmetrical and more restricted during phonation. This paper proposes a method that analyzes the differential dynamics of the glottal flow waveform (GFW) during voice production to identify features in them that are most significant for the detection of COVID-19 from voice. Since it is hard to measure this directly in COVID-19 patients, we infer it from recorded speech signals and compare it to the GFW computed from physical model of phonation. For normal voices, the difference between the two should be minimal, since physical models are constructed to explain phonation under assumptions of normalcy. Greater differences implicate anomalies in the bio-physical factors that contribute to the correctness of the physical model, revealing their significance indirectly. Our proposed method uses a CNN-based 2-step attention model that locates anomalies in time-feature space in the difference of the two GFWs, allowing us to infer their potential as discriminative features for classification. The viability of this method is demonstrated using a clinically curated dataset of COVID-19 positive and negative subjects.

**Index Terms**— COVID-19 detection, Vocal fold oscillation, Models of phonation, Attention models, Interpretable neural nets

## 1. INTRODUCTION

The production of voiced sounds is a two-step process involving the self-sustained vibrations of the vocal folds, and the propagation of the resultant pressure wave within the vocal tract [1]. In this paper, we focus on the first step: the motion of the vocal folds. Recent studies have shown that this motion is adversely affected in symptomatic COVID-19 patients [2]. While these studies are able to identify broad-level anomalies (such as asynchrony, asymmetry of motion, reduced range of oscillation etc.) by visual comparisons between oscillation patterns of healthy and *symptomatic* COVID-19 positive people, more systematic methods of comparison are needed to identify discriminative signatures that may be used as features

in classifiers that aim to detect COVID-19 from voice.

This problem may be approached by first noting that the motion of the vocal folds is highly correlated with the size of the glottal opening. The glottis opens and closes as the vocal folds vibrate. In this self-sustained cycle of the opening and closing of the glottis, air pressure periodically builds up and is released, giving rise to a fluctuating volume of airflow across the glottis that comprises the glottal flow waveform (GFW). Since the glottal opening and closing is caused by the motion of the vocal folds, the spatial displacements of the vocal folds are normally highly correlated to the GFW. Simple physical models of phonation express these displacements as a function of the bio-physical properties of the vocal folds, such as their thickness, elasticity etc. and of the aerodynamic forces across the glottis e.g. [3, 4, 5, 6, 7]. Given a set of parameters, these models yield the oscillation of the vocal folds.

Ideally, if we were to compute the difference in the GFW of the same speaker when healthy and when symptomatic with COVID-19, we would learn a lot from the differences we observe. However, such pairs of recordings are not easily available. In lieu of this, if only a recording of a COVID-19 positive individual is available, we can compute the GFW (denoted  $GFW_{filter}$ ) from it by inverse filtering the speech signal. In this work we assume that we do *not* have a recording of the same speaker when healthy. We now *deliberately* assume that a physical model that explains the vocal fold motions of a healthy individual will match the those of the COVID-19 positive speaker obtained from inverse filtering. Clearly, this assumption is not true, and will be invalidated to different degrees, depending on how much it deviates from reality. Due to the correlation of the GFW to the vocal fold oscillations, these deviations are expected to surface in the form of differences between the GFW calculated from the estimated vocal fold motions from the physical models ( $GFW_{model}$ ) and  $GFW_{filter}$ . These differences can then be ascribed to the influence of COVID-19, which changes the underlying bio-physical factors that the model attempts to emulate, and renders the assumptions invalid in proportion to the severity of the disease in the speaker.

In this paper, we the ADLES [8] algorithm to jointly estimate the vocal fold oscillations and  $GFW_{model}$  and compare it to the  $GFW_{filter}$  obtained from the speech signal. We propose a method of analysis of the observed differences,

to identify the features in it that are relevant to classifiers for COVID-19 detection. The method uses an interpretable CNN model with a 2-step attention pooling architecture to detect the differences (residual) as a function of time. We show that it successfully identifies the traces that are of significance to COVID-19 detection from voice, and leads to better classification performance when tested on a clinically curated dataset of COVID-19 positive and negative individuals.

## 1.1. Related Work

The literature on detecting COVID-19 from voice and respiratory sounds is relatively new [9] and sparse. Literature on using voice, and especially the dynamics of vocal fold motion to detect COVID-19 is even more sparse, with only one paper in existence [2]. To the best of our knowledge, our current paper is the first one that addresses the problem of interpretation of vocal flow dynamics for more accurate identification of the discriminative factors that can be used for classification of COVID-19 from voice.

## 2. PROPOSED METHODOLOGY

### 2.1. Computing the GFW from a phonation model

We use the 1-mass asymmetric body cover model proposed in [10], and using the ADLES algorithm proposed in [8], compute oscillation of the vocal folds from which an estimate for the glottal flow waveform  $GFW_{model}$  can be obtained. To compute the vocal fold oscillations (or the displacements of the vocal folds), the ADLES algorithm solves the *forward* problem of estimating the displacement and velocity of each vocal fold jointly with the *inverse* problem of estimating the parameters of the dynamical system comprising the model. To do this, we must minimize the squared difference between the estimated glottal flow from inverse filtering  $GFW_{filter}$ , and  $GFW_{model}$ . Mathematically:

$$\min \int_0^T (u_0(t) - u_0^m(t))^2 dt \quad (1)$$

$$\min \int_0^T (\tilde{c}d(2x_0 + x_l(t) + x_r(t)) - \frac{A(0)}{\rho c} \mathcal{F}^{-1}(p_m(t)))^2 dt \quad (2)$$

$$\text{s.t. } \dot{x}_r + \beta(1 + x_r^2)\dot{x}_r + x_r - \frac{\Delta}{2}x_r = \alpha(\dot{x}_r + \dot{x}_l) \quad (3)$$

$$\dot{x}_l + \beta(1 + x_l^2)\dot{x}_l + x_l + \frac{\Delta}{2}x_l = \alpha(\dot{x}_r + \dot{x}_l) \quad (4)$$

$$x_r(0) = C_r, x_l(0) = C_l, \dot{x}_r(0) = 0, \dot{x}_l(0) = 0 \quad (5)$$

where the Eqns. (3) and (4) jointly correspond to the model.  $x_r, x_l$  are the displacements of right and left vocal fold,  $\dot{x}_r, \dot{x}_l$  are horizontal vocal fold velocity of right and left vocal fold w.r.t the center of the glottis,  $\alpha, \beta$  and  $\Delta$  are the model's parameters.  $\alpha$  is the coupling coefficient between the supra-

and sub-glottal pressure,  $\beta$  incorporates the mass, spring and damping coefficients of the vocal folds, and  $\Delta$  is an asymmetry coefficient. Here  $GFW_{filter} \equiv u_0^m(t)$  is the GFW obtained from inverse filtering, and  $GFW_{model} \equiv u_0(t)$  is estimated GFW as a joint solution to the model with ADLES.

Within ADLES, the optimization problem is cast into its dual by using Lagrangian followed by optimization using gradient descent. The process of obtaining the solution for parameters is explained in detail here [11].

### 2.2. COVID-19 analysis framework

We represent each recording as a sequence of sliding windows (frames)  $\{X_i\}_{i=1}^T$ , where each frame is of duration  $\tau$ , and the frame shift is  $o$ . For each of the frames  $X_i$ , a  $GFW_{filter} \equiv u_i(t)$  estimate is obtained by inverse filtering and also as a function of vocal fold motion  $GFW_{model} \equiv u_i^m(t)$ . We will drop the subscript 0 indicating measurement at the glottis from this notation for brevity. Each frame is thus represented by  $[u_i(t), u_i^m(t)]$ , and the sequence of audio frames is now  $\{X_i\}_{i=1}^T = \{[u_i(t), u_i^m(t)]\}_{i=1}^T \in \mathbb{R}^{2 \times T}$ .

We wish to capture the differences between the two GFW estimates in a way that a classifier is best able to use them to discriminate between COVID-19 positive and negative cases, i.e., we want to estimate the probability  $P(y|u_i(t), u_i^m(t); \mathbf{w})$  where  $\mathbf{w}$  represents the parameters of a parametric model.  $y \in \{0, 1\}$  is a set of binary labels used with a classifier. The task of estimating  $P(y|u_i(t), u_i^m(t); \mathbf{w})$  can be divided into two stages: in the first, a time-distributed pattern detector is trained to find the similarities and differences between  $u_i(t), u_i^m(t)$  at each time step (a frame comprises  $N$  time steps). In the second, a pooling mechanism aggregates the outputs of the first stage to yield a single prediction for each frame.

Let the output of first stage be given by  $\mathbf{Z}$  and let the learnable feature extractors be represented by  $f_1(\cdot)$ . The goal of  $f_1(\cdot)$  is to learn a mapping such that:

$$f_1 : \{[u_i(t), u_i^m(t)]\} \mapsto \mathbf{Z}$$

The aim of the second stage is to reduce  $\mathbf{Z}$  from stage 1 along both, the time and feature axes to obtain single frame-level prediction  $y_i$ . This can be thought of as a multiple-instance learning (MIL) setup [12] where each window/frame  $X_i$  is represented as a bag of  $N$  samples:  $B_i = (\{x_j\}_{j=1}^N, y_i)$  where  $x_i$  represents individual samples in the frame. This type of MIL formulation has been extensively used in a Weakly Labeled Sound Event Detection (WLSSED) [13] framework. The pooling methods developed for WLSSED [14, 15, 16] are also transferable to our problem of MIL based COVID-19 detection. We now introduce a pooling function  $f_2(\cdot)$  that learns the mapping:  $f_2 : \mathbf{Z} \mapsto y$

The 2-step Attention Pooling (2AP) mechanism [16] yields benchmark performance for WLSSED, in the task of determining the contribution of each class at each instance in

both time and frequency domains. In the problem setup in [16], the output of the first step of attention is  $C \times T$  where  $C$  is the number of classes and is the result of a weighted combination of the frequency components. In our case, as we have a single class setup (COVID/non-COVID), the output of first step of attention is  $1 \times T$ , which is a weighted combination of the amplitudes of  $GFW_{filter}$  and  $GFW_{model}$ . This weighted combination of amplitudes ranges from high to low depending on whether the phase difference between the two signals is low or high respectively. Adding an extra feature detector  $f_3(\cdot)$  between the two stages allows the model to better capture the entities of interest to us as a function of time. We call this step the Sandwiched 2-step Attention Pooling (S2AP) step. Subsequently, we use S2AP for learning the mapping  $f_2(\cdot)$ . Formally, first step in S2AP takes  $\mathbf{Z}$  as input:

$$\hat{Z}_{a1} = \frac{e^{\sigma(\mathbf{Z}\mathbf{w}_{a1}^T + b_{a1})}}{\sum_{i=1}^F e^{\sigma(\mathbf{Z}\mathbf{w}_{a1}^T + b_{a1})}}, Z_{p1} = \sum_{i=0}^F (\mathbf{Z}\mathbf{w}_{c1}^T + b_{c1}) \cdot \hat{Z}_{a1} \quad (6)$$

$$Z_x = f_3(Z_{p1}) \quad (7)$$

$$\hat{Z}_{a2} = \frac{e^{\sigma(Z_x \mathbf{w}_{a2}^T + b_{a2})}}{\sum_{t=1}^T e^{\sigma(Z_x \mathbf{w}_{a2}^T + b_{a2})}}, Z_{p2} = \sum_{t=0}^T (Z_x \mathbf{w}_{c2}^T + b_{c2}) \cdot \hat{Z}_{a2} \quad (8)$$

where  $\sigma(\cdot)$  is Sigmoid function,  $\{\mathbf{w}_{a1}, \mathbf{w}_{c1}, \mathbf{w}_{a2}, \mathbf{w}_{c2}\}$  are learnable weight parameters and  $Z_{p2}$  is  $y$  and gives the probability of the voice sample containing traces of COVID-19. Moreover, S2AP retains the benefits of 2AP of making the pooling more interpretable. Visualization of the normalized attention weights  $\hat{Z}_{a1}$ ,  $\hat{Z}_{a2}$  and the output  $Z_{p1}$ ,  $Z_{p2}$  allows us to infer what features and what time steps contribute most significantly to the decisions made by the classifier.

### 3. EXPERIMENTS

**Dataset:** We used recordings of extended vowels /a/, /i/, /u/ of COVID-19 positive and negative (medically tested) individuals. The data were collected under clinical supervision and curated by Merlin Inc., a private firm in Chile. Of 512 data instances, we chose 19 recordings that were collected within 7 days of testing. These comprised 10 females (5 positive) and 9 (4 positive) males. The recordings comprised 8kHz sampled signals recorded on commodity devices.

**Setup:** Fig. 1 describes the overall setup used for detection of COVID-19. For each file we use a window  $\tau = 50ms$  with shift  $o = 25ms$ , resulting in 3835 frames in all. The extracted  $GFW_{filter}$  and  $GFW_{model}$  were analyzed with a CNN classifier for the a binary COVID-19 detection task. 3-fold cross validation experiments were performed, where the folds were stratified by speaker: no speaker in the training set was included in the test set, to ensure that the classifier did not learn speaker dependent characteristics. The performance metrics used were the area under curve (ROC-AUC) of the the receiver operating characteristics graph, and its standard

deviation (STD).  $GFW_{filter}$  was obtained using the IAIF algorithm [17] and  $GFW_{model}$  was obtained using the ADLES algorithm.

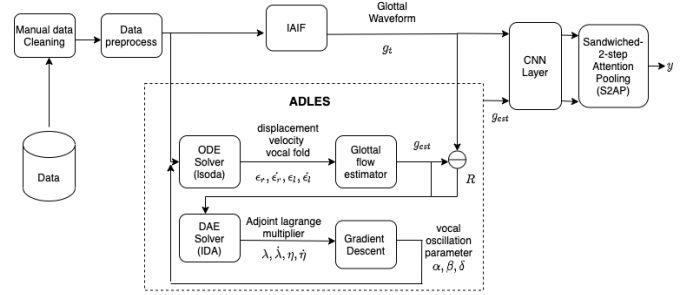


Fig. 1. Covid-19 analysis system setup

### 4. RESULTS

**Classifier results:** Table 1 shows the classification results in terms of ROC-AUC (averaged) and STD in the 3-fold cross-validation experiment. The classifier used was a convolutional neural network (CNN).

Feature extractor	Pooling	ROC-AUC	STD
-	2AP	0.6611	0.0978
-	2SAP	0.7925	0.1073
CNN (1,3,32)	2AP	0.8009	0.1009
CNN (2,3,32)	2AP	0.8248	0.0790
CNN (2,3,32)	2SAP	0.8330	0.0745
CNN (2,5,64)	2SAP	<b>0.8520</b>	<b>0.0577</b>

Table 1. Classifier performance on 3-fold cross validation

	/a/	/i/	/u/	/a+/i/	/a+/u/	/i+/u/
AUC	0.57	0.839	0.896	0.690	0.804	<b>0.900</b>
STD	0.119	0.102	0.067	0.064	0.074	<b>0.062</b>

Table 2. Performance of best model on individual extended vowels and their combinations

We initially use no feature extractors, and directly do a 2-step attention pooling (2AP). This achieves an ROC-AUC of 0.66, indicating the presence of large differences (lack of one-to-one correspondence) between  $GFW_{filter}$  and  $GFW_{model}$ . This informs us that neighboring time instances must also be considered. We therefore introduce a new CNN layer with a kernel size of 3 to capture the immediate neighborhoods samples in the frame. In the table 1, the label CNN (1,3,32) indicates 1 CNN layer with kernel size 3 and filter size 32. This improves the ROC-AUC by 21% over the best performance without this layer.

Visualizing the attention patterns reveals that this layer highlights the phase difference and residual between  $GFW_{filter}$

and  $GFW_{model}$ . To better detect the patterns temporally (along the time-axis), we then introduce a second CNN layer to increase the receptive field (denoted by 2,3,32/2AP in Table 1). This is known to help capture *long term* patterns. This results in a further 1% relative improvement.

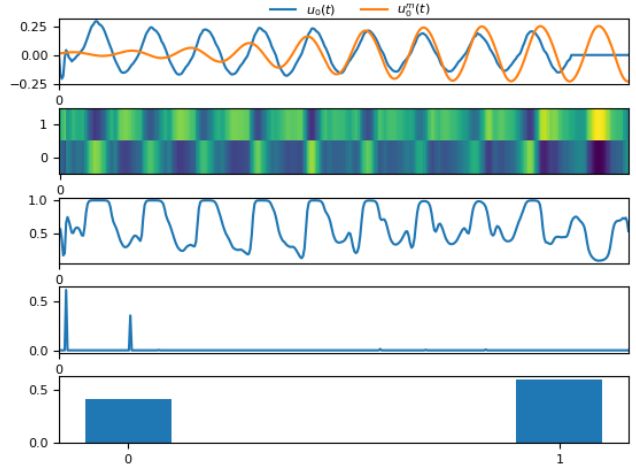
As explained in Sec. 2.2, we now introduce 2SAP (denoted by 2,3,32/2SAP in Table 1). This improves performance by 19.8% compared to the standard 2AP, achieving an ROC-AUC of 0.833 when coupled with the initial feature extractor. The final model (kernel size 5, no. of filters 64, denoted by 2,5,64/2SAP) further increases the receptive field [18] and achieves our benchmark ROC-AUC of 0.852.

Note that the model used here is relatively simple compared to the deep stacked CNNs used in other speech and audio classification tasks. Deep-stacking the CNN layers may improve the performance further, but CNN architectures are not focus of this paper. By using a couple of simple CNN layers here, we show the utility of the proposed methodology in identifying the features that best capture the anomalies in vocal fold motion in a manner that optimizes classifier performance for the current task (detection of COVID-19).

**Vowel analysis:** Table 2 shows the results obtained when our best model derived above (CNN(2,3,64)/2SAP) is trained on particular vowels. The best performance achieved (ROC-AUC of 0.9) is for the combination of vowels /u/ (high back vowel) and /i/ (high front vowel). The vowel /a/ (low back vowel) exhibits the least importance in this task. This seems to indicate that the vocal folds are unable to vibrate such that higher harmonics have normal energy distributions. However, this conjecture remains to be validated using actual measurements of vocal fold movements.

**Interpretability:** A key advantage of using S2AP is its ability to visualize the contribution of each feature and each time step to classifier’s task (the binary task of detecting the presence of COVID-19 in this case). We illustrate this with the example in Fig. 2, that shows the attention analysis of a randomly selected frame (a 50ms window) from a COVID-19 positive patient. Going from top to bottom in this figure, panel 1 shows the feature streams:  $GFW_{filter}$  and  $GFW_{model}$  for this frame. Initially,  $GFW_{model}$  is not able to match  $GFW_{filter}$ . Panel 2 shows the attention weights from the 1<sup>st</sup> step of the proposed method. The features are on  $y$ -axis and time is on the  $x$ -axis. The feature stream labeled 0 corresponds to  $GFW_{model}$  and the one labeled 1 corresponds to  $GFW_{filter}$  estimated using IAIF. As we see, the 1<sup>st</sup> step attention alternately focuses on feature 0 and 1 to detect whether the two features match or not, i.e. on the residual as a function of time. In this panel the lighter shades indicate higher values of the attention weights.

The output of 1<sup>st</sup> step attention is visualized in the third panel. It is high when the peaks and lows of two waves are in synchrony, and can be thought of as detecting phase-locking of the two waves, weighted by the attention. The fourth panel shows the time-level (2<sup>nd</sup> step) attention weights.



**Fig. 2.** Visualization of attention weights and in the Sandwiched 2-step attention pooling scheme

The 2<sup>nd</sup> step attention output peaks at the initial point where two GFWs not synchronized. This is the time that contributes the most significant information for the classification task. The fifth panel is a histogram with 2 bins, centered at 0 and 1. This corresponds to the output of the binary classifier obtained after 2<sup>nd</sup> step attention. The classifier gives higher weight to ( $> 0.5$ ) to class 1, which indicates the presence of COVID-19 correctly in the chosen frame, as confirmed by its true label.

Being able to visualize what the classifier focuses in this task paves the way for further improving it by examining its mistakes. It also provides human interpretability – a way of adding expert human intervention to black-box model decisions. For example, if the model uses an erroneous pattern leading to decision of COVID-19 traces in the data, it can be discarded as a feature if the human expert does not agree with what entity the attention is focusing on.

## 5. CONCLUSIONS

For the detection of COVID-19 from voice, we introduce a method for identifying meaningful features from a comparison of the GFW from inverse filtering to that obtained by solving a dynamical system model for vocal fold oscillations. Our method uses a novel attention mechanism (Sandwiched 2-step Attention Pooling (S2AP)) built within a CNN framework to identify the most discriminative features. For this task it shows that the residual and the phase difference between the two GFWs are the most promising features. Using these for COVID-19 detection yields an ROC-AUC of 0.9 when used with the extended vowels /u/ + /i/. Overall, this methodology paves the way for comparing any set of feature streams for feature discovery, and can potentially be used in any general classification task where the comparison of two feature streams is expected to yield meaningful features.

## 6. REFERENCES

- [1] Ingo R Titze, “Nonlinear source–filter coupling in phonation: Theory,” *The Journal of the Acoustical Society of America*, vol. 123, no. 4, pp. 1902–1915, 2008.
- [2] Mahmoud Al Ismail, Soham Deshmukh, and Rita Singh, “Detection of covid-19 through the analysis of vocal fold oscillations,” 2020.
- [3] Jorge C Lucero and Jean Schoentgen, “Modeling vocal fold asymmetries with coupled van der pol oscillators,” in *Proceedings of Meetings on Acoustics ICA2013*. ASA, 2013, vol. 19, p. 060165.
- [4] Kenzo Ishizaka and James L Flanagan, “Synthesis of voiced sounds from a two-mass model of the vocal cords,” *Bell system technical journal*, vol. 51, no. 6, pp. 1233–1268, 1972.
- [5] Anxiong Yang, Michael Stingl, David A Berry, Jörg Lohscheller, Daniel Voigt, Ulrich Eysholdt, and Michael Döllinger, “Computation of physiological human vocal fold parameters by mathematical optimization of a biomechanical model,” *The Journal of the Acoustical Society of America*, vol. 130, no. 2, pp. 948–964, 2011.
- [6] Fariborz Alipour, David A Berry, and Ingo R Titze, “A finite-element model of vocal-fold vibration,” *The Journal of the Acoustical Society of America*, vol. 108, no. 6, pp. 3003–3012, 2000.
- [7] Ingo R Titze, “The physics of small-amplitude oscillation of the vocal folds,” *The Journal of the Acoustical Society of America*, vol. 83, no. 4, pp. 1536–1552, 1988.
- [8] W. Zhao and R. Singh, “Speech-based parameter estimation of an asymmetric vocal fold oscillation model and its application in discriminating vocal fold pathologies,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7344–7348.
- [9] Gauri Deshpande and Björn Schuller, “An overview on audio, signal, speech, & language processing for covid-19,” 2020.
- [10] Jorge C. Lucero, Jean Schoentgen, Jessy Haas, Paul Luizard, and Xavier Pelorson, “Self-entrainment of the right and left vocal fold oscillators,” *The Journal of the Acoustical Society of America*, vol. 137, no. 4, pp. 2036–2046, 2015.
- [11] Wenbo Zhao and Rita Singh, “Speech-based parameter estimation of an asymmetric vocal fold oscillation model and its application in discriminating vocal fold pathologies,” 2020.
- [12] Thomas G. Dietterich, Richard H. Lathrop, and Tomás Lozano-Pérez, “Solving the multiple instance problem with axis-parallel rectangles,” *Artif. Intell.*, vol. 89, no. 1–2, pp. 31–71, Jan. 1997.
- [13] Anurag Kumar and Bhiksha Raj, “Audio event detection using weakly labeled data,” in *Proceedings of the 24th ACM International Conference on Multimedia*, New York, NY, USA, 2016, MM ’16, p. 1038–1047, Association for Computing Machinery.
- [14] Qiuqiang Kong, Yong Xu, Iwona Sobieraj, Wenwu Wang, and Mark D. Plumbley, “Sound event detection and time–frequency segmentation from weakly labelled data,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 27, no. 4, pp. 777–787, Apr. 2019.
- [15] Z. Ren, Q. Kong, J. Han, M. D. Plumbley, and B. W. Schuller, “Attention-based atrous convolutional neural networks: Visualisation and understanding perspectives of acoustic scenes,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 56–60.
- [16] Soham Deshmukh, Bhiksha Raj, and Rita Singh, “Multi-task learning for interpretable weakly labelled sound event detection,” 2020.
- [17] Paavo Alku, “Glottal wave analysis with pitch synchronous iterative adaptive inverse filtering,” *Speech communication*, vol. 11, no. 2-3, pp. 109–118, 1992.
- [18] Hendrik Purwins, Bo Li, Tuomas Virtanen, Jan Schluter, Shuo-Yiin Chang, and Tara Sainath, “Deep learning for audio signal processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 206–219, May 2019.