DETECTION OF SPEECH EVENTS AND SPEAKER CHARACTERISTICS THROUGH PHOTO-PLETHYSMOGRAPHIC SIGNAL NEURAL PROCESSING

Guillermo Cámara^{1,2}, Jordi Luque^{1,3} and Mireia Farrús²

¹Telefonica Research, Barcelona, Spain

²Universitat Pompeu Fabra, Barcelona, Spain

³Universitat Politècnica de Catalunya, Barcelona, Spain

ABSTRACT

The use of photoplethysmogram signal (PPG) for heart and sleep monitoring is commonly found nowadays in smart-phones and wrist wearables. Besides common usages, it has been proposed and reported that person information can be extracted from PPG for other uses, like biometry tasks. In this work, we explore several end-to-end convolutional neural network architectures for detection of human's characteristics such as gender or person identity. In addition, we evaluate whether speech/non-speech events may be inferred from PPG signal, where speech might translate in fluctuations into the pulse signal. The obtained results are promising and clearly show the potential of fully end-to-end topologies for automatic extraction of meaningful biomarkers, even from a noisy signal sampled by a low-cost PPG sensor. The AUCs for best architectures put forward PPG wave as biological discriminant, reaching 79% and 89.0%, respectively for gender and person verification tasks. Furthermore, speech detection experiments reporting AUCs around 69% encourage us for further exploration about the feasibility of PPG for speech processing tasks.

Index Terms— Photoplethysmogram signal, PPG, speech detection, convolutional neural networks, biometric authentication

1. INTRODUCTION

Nowadays, tens of biometric sensors are commonly embedded in most of electronic personal devices like laptops, smart-phones or smart-watches. Typically, these sensors have been used to retrieve information like heart rate, blood oxygen level or user's fingerprint, for applications ranging from health-care to biometric-based authentication. Previous studies [1, 2, 3] have found that biometric signals obtained from sensors such as gyroscopes, accelerometers, or photoplethysmograms (PPG) can be processed to extract both further information and different patterns from those they were originally designed for. Furthermore and together with the advent of end-to-end neural network architectures, showing astonishing performances, they might translate into novel use cases

and applications for biometrics in common wearable devices. Some recent works have studied the feasibility of using the PPG signal [4] for user authentication, as shown in [2], [3]. Typically, these methods involve both expert knowledge and considerable effort for the extraction of relevant biomarkers, which can encode user's characteristics. Recently, in [1], it was shown that Convolutional Neural Network (CNN) methods can be used for automatically extracting such relevant features, proposing an end-to-end neural network architecture for the task of user's identity verification. The proposed fully end-to-end system, capable of automatically extract discriminant biomarkers, outperforms classical reported methods in both TROIKA [5] and self-recorded databases.

Such findings motivate this work, which intends to extract relevant information from PPG signal and open up a door for further application possibilities, besides other recent proposed applications as user authentication. Some works have proven that speech information can be mined using indirect devices other than microphones, like a gyroscope [6] or an accelerometer [7]. So, a PPG sensor could be a proper candidate for detecting speech events, as any other sensor present in smart devices.

We hypothesize that speech produced by the user causes some sort of fluctuation in the PPG signal, either directly by low-frequency acoustic components from speech itself, or indirectly by the sensor movement [6, 7]. Some authors also have pointed that speech events might translate into PPG by an additional oxygen consumption [8]. In any case, the detection of such fluctuations could suggest the existence of a speech event during a certain period of time. Wrapping up, the main aim of this work is to discern whether there is any correlation between speech/non-speech events and PPG. Furthermore, our aim is to explore the feasibility to distinguish between different speech events as PPG serves to discriminate between users by detecting gender or identity. The latter would allow to develop tasks like speech/non-speech detection, speech enhancement or even word recognition, e.g., by augmenting current methods with PPG information.

Fig. 1: Multimodal 8-second recording of speaker 004 from PulseID database. On top, spoken digits sequence, in the middle the audio waveform and, at the bottom the PPG signal. Dashed vertical lines depict the forced-alignment obtained between text and audio – pronounced digits – given by an automatic speech recognition system

2. PREVIOUS WORK

Up to the authors’ knowledge, no previous work has tried to find the correlation between PPG signal and speech, with or without deep learning methods. So, we will focus on the literature related to the topics converging in this work, such as speaker identification by using PPG signal, speech detection with unconventional devices, oxygen consumption on speech production or deep neural networks (DNN) applied to noise detection in electrocardiogram signal (ECG).

Some previous works [2, 3] studied time domain characteristics like time intervals, peaks, upward and downward slopes in PPG signal, used to infer a person’s identity. In [1] it was shown that a CNN architecture allows for direct classification without feature extraction effort. Further related works [6] found out that gyroscopes capture acoustic vibrations from speech under frequencies of 200 Hz, which can be used to detect speech, identify speakers, or even parse such speech, using Short Time Fourier Transform (STFT), Mel-Frequency Cepstral Coefficients as features and by using machine learning classification methods like Support Vector Machine (SVM), Gaussian Mixture Model or Dynamic Time Warping. In a similar way, [7] reported that acoustic vibrations leaked into an accelerometer signal as well, thus allowing for speech activity detection. They performed feature extraction using the Fast Fourier Transform (FFT), and various machine learning algorithms such as SVM or Naïve Bayes. Therefore, low frequencies from speech might leak as well into a PPG sensor, opening up for eavesdropping possibilities.

Similarly, noise detection task was done in [9], achieving the best results with a 16-layer CNN adapted from the

VGG16 network [10]. However, besides acoustic vibrations, an interesting indicator of speech might be the difference in oxygen consumption between speech and rest, measured by the PPG signal. In [8] it was proven that oxygen consumption increases when the vocal effort (speech volume) is higher, and that a higher frequency of pronounced syllables yields a higher oxygen consumption as well. The first result reinforces the hypothesis that speech or non-speech events could be classified from a PPG signal because of the oxygen concentration in blood.

3. METHODOLOGY

3.1. Data source synchronization

The PulseID dataset introduced in [1] is employed to assess the proposed CNN architectures. It was collected using the open source and low-cost photoplethysmograph sensor described in [11], which consists of a green LED and a photo-detector, sampling at a 200 Hz rate. For a sample of PPG signal, see Figure 1.

PPG waveform is synchronized up with audio recordings at same PPG recording device. For this purpose, a developed python code exploits threading capabilities of a Raspberry Pi 3 device, see [1, 12] for more details. The analysis of error measurements shows a tolerant sampling rate deviation ($\mu = 13.32 \mu s$ and $\sigma = 202.58 \mu s$) per subject, which lies within an acceptable margin for speech events. Audio signal is recorded at a 44.1 kHz sample rate and 16 bits per sample, but a downsampling to 8 kHz is done for further processing.

3.2. Recording protocol

Data acquisition implied the participation of 31 subjects (25 males and 6 females), with ages ranging from 22 to 55 years old. Five different types of speech recordings were performed through a microphone, including the pronunciation of certain words and sequence numbers like PINs. All the experiments include at least 30 seconds of pulse and audio signals, except for the fifth one, which took place during 1 minute.

The first experiment involves subjects saying two random credit card numbers of 16 digits each at a regular pace, with a longer pause between both numbers. The second one contains a 30-second recording of the pulse and the audio without any speech from the subject. The third experiment, similar to the first one, includes subjects uttering four random 6-digit PIN numbers, also with a longer pause between two consecutive PIN numbers. The fourth experiment uses phonetically rich sentences commonly employed for ASR corpus, uttered by the subjects at a regular pace and with a longer pause between sentences. The last experiment includes one minute of free speech, where the subjects typically describe their environment¹. Both phonetically rich sentences and protocol guidelines are following those in [13].

3.3. Audio and text forced-alignment

Once the PPG and the audio signals are synchronously recorded, the latter is used for getting speech timestamps at the frame level. The speech labelling is performed by forced-alignment of text sentences/numbers and audio. It is accomplished by an Automatic Speech Recognizer (ASR), previously trained with more than two hundred hours of Spanish speech and the Kaldi toolkit [14]. We make use of one of prototypes in [15]. Since the words pronounced are known in advance – except for the free speech case, which is not annotated — the ASR performs a lattice decoding and iteratively refines it by seeking the phoneme sequence that maximizes a given word sequence. Nonetheless, previous methodology produces few errors, especially whenever strong background noise happens, so still a manual cleaning is done and non-accurate time stamps are corrected by hand.

3.4. Neural network architectures

Five different CNN based architectures are benchmarked for speech detection, speaker identification and gender classification tasks: the net from [1] – from now on PulseNet – plus two variants, the inverted VGG-like model employed for noise detection in ECG signals [9] and a CNN which uses 2D STFT as an input – CNN-2D –, adapted from a kernel in Kaggle’s FAT 2019 competition [16]. Standard models like GMM-UBM or i-vector are not used due to the small

¹The *PulseID* dataset, including audio and alignments, is available upon request from the authors and agreement of EULA for research purposes.

Fig. 2: Boxplot of AUC values for test set evaluation given by four different architectures. Results are depicted across 100 repetitions

size of the dataset, which would not provide a good estimation for the background model (GMM-UBM) nor the matrix T (i-vectors). The PulseNet initial kernel sizes ($L_1, L_2, L_3 = 50, 30, 20$) for the three convolutional layers reported in [1] are optimized by brute force, exploring configurations of kernel sizes $L_i \in [2, 3, 4, \dots, 180, 200]$. Using a PulseNet model with smaller kernel sizes aims at exploring shorter time correlations. From such explorations, we report the results on the best performing variants: PulseNet-Var1 ($L_1, L_2, L_3 = 50, 10, 4$) and PulseNet-Var2 ($L_1, L_2, L_3 = 15, 8, 2$).

Cross-entropy loss function is used together with a variant of SGD optimizer, known as SGD with Restarts [17]. It consists of decreasing the learning rate in the form of half a cosine curve, restarting it at the minimum to the original value. Batch size and learning rate are fine tuned for every architecture, allowing the model to safely decrease the loss whilst augmenting label accuracy. The code to replicate following experiments, mainly written using PyTorch [18] and initially forked from the implementation in [19], can be found at [20].

4. EXPERIMENTS AND RESULTS

For every PPG recording, the signal is split with a rolling window, which is sliced along the time axes. Every PPG window sample is marked as *speech* or *non-speech*, as indicated by audio timestamps (see section 3.3). Since there are short moments during speech where the speaker stops and breathes, if a *speech* signal window contains more than a 2% of silence, it is discarded, ensuring that the speaker was speaking the most of the time during signal samples labeled as *speech*. After exploring several sizes and strides for the rolling window, the 1-second window as in [1] seems to yield the best results. Some experiments are done subdividing such 1-second win-

Arc	Window Size (s)	Speech Detection		Gender Detection		Speaker Verification	
		AUC	F1-Score	AUC	F1-Score	AUC	F1-Score
PulseNet-Var1	1	68.2 $\pm$ 0.3	63.9 $\pm$ 0.3	77 $\pm$ 5	61 $\pm$ 5	80.0 $\pm$ 0.1	60.7 $\pm$ 0.2
PulseNet-Var2	1	67.6 $\pm$ 0.3	64.0 $\pm$ 0.3	75 $\pm$ 5	60 $\pm$ 5	74.4 $\pm$ 0.1	48.7 $\pm$ 0.3
PulseNet	1	66.8 $\pm$ 0.3	63.3 $\pm$ 0.3	74 $\pm$ 2	60 $\pm$ 2	82.7 $\pm$ 0.1	66.1 $\pm$ 0.3
VGG16-Inv	1	64.6 $\pm$ 0.3	62.6 $\pm$ 0.3	68 $\pm$ 2	55 $\pm$ 2	83.4 $\pm$ 0.1	67.3 $\pm$ 0.2
CNN-2D	1	61.7 $\pm$ 0.3	60.4 $\pm$ 0.3	79 $\pm$ 1	66 $\pm$ 1	89.0 $\pm$ 0.1	78.5 $\pm$ 0.2
PulseNet-Var1	0.4	66.7 $\pm$ 0.3	64.0 $\pm$ 0.3	70 $\pm$ 2	59 $\pm$ 2	68.4 $\pm$ 0.1	37.9 $\pm$ 0.1
PulseNet	0.4	67.6 $\pm$ 0.3	64.7 $\pm$ 0.3	66 $\pm$ 2	56 $\pm$ 2	70.7 $\pm$ 0.1	42.3 $\pm$ 0.1

Table 1: Result scores for speech detection, gender detection and speaker verification test tasks, averaged over 100 repetitions, errors given by the standard error of the mean. PPG signals are split in smaller samples according to the window size column and fed to five different CNN architectures.

dows in smaller windows of 0.4 seconds, each one of them overlapped in a 70% with the adjacent ones, to explore the effect of overlapping during training. The decision on the final label of a 1 second sample in the test is done by summing the log-probabilities obtained from the forward pass of each 0.4-second sub-sample.

Every result score is averaged along 100 experiments, each one of them with shuffled training (64%), validation (16%) and testing (20%) partitions. However, for the gender classification experiment the samples from 2 males and 2 females speakers are kept exclusively only for the testing set. 20 folds with different male-female combinations are tested, averaging the performance scores through 25 experiment repetitions with shuffled training and validation sets. Because of the bigger size in the CNN-2D network, the repetitions have been reduced to 10 times for a 10-fold in the gender detection task and 22 times for the speaker verification one, in order to save computational time. AUC and F1 weighted scores measure how good the model is, preferred over accuracy, in order to account for class imbalances, specially for speech detection (33%/66% for speech/non-speech) and gender classification (81%/19% for male/female).

As can be seen in Table 1 and Figure 2, the best results for speech detection are achieved by the PulseNet architectures, attaining a highest scoring mean AUC of $68.2 \pm 0.3\%$ with PulseNet [50,10,4]. This configuration seems to benefit of smaller kernel sizes than the original PulseNet to find correlations in the signal, suggesting that speech fluctuations could be really fine-grained. Using the original PulseNet with a smaller 0.4 s window sizes practically matches this result without need of reducing the kernels, which also points out in the direction of zooming in in the signal to find such correlations. Despite the CNN-2D architecture yields the worst result for speech detection, it achieves the best scores for gender detection and speaker verification tasks, with $79 \pm 1\%$

and $89.0 \pm 0.1\%$ AUCs respectively, surpassing PulseNet’s performance particularly for the latter. The VGG16-Inv model shows fairly good performance at all tasks, but does not achieve any highest scoring result. It seems that mining speaker characteristics is better done with deeper CNNs with varying receptive fields at each filter, which is the case for CNN-2D and VGG16 architectures. Particularly, an approach with a STFT spectrogram like in CNN-2D might be better than using just the raw signal, as opposed to the speech detection task, which benefits from the smaller and parallel filters of PulseNet architectures. High variance in gender results is due to the low number of subjects, specially females. Depending on the speaker, the network achieves excellent or poor detection rates.

5. CONCLUSION

In this work, we have explored end-to-end CNN architectures to detect speech and users’ gender and identity characteristics. We have found that all five architectures used are able to distinguish speech from silence through PPG signals, not with a high accuracy, but indeed way over the random baseline, with all AUC and F1 scores over 60%, averaged in runs of 100 experiments, and achieving almost a 70% AUC with the best architecture. Such model is a variant of the original PulseNet configuration, with smaller kernel sizes, which suggests that speech information might appear in PPG signal as some sort of fine-grained fluctuation. Such results motivate further research, in order to clarify the nature of these fluctuations, and to improve the performance of existing systems, which would benefit of larger datasets and less noisy sensors. Furthermore, bi-dimensional CNN architectures using STFT spectrograms have shown to be the best option for extracting speaker characteristics like gender or identity, improving the performance of the PulseNet network for the latter task.

6. REFERENCES

- [1] Jordi Luque, Carlos Segura, Guillem Cortes, Alexandre Maravilla, Javier Esteban Zarza, and Joan Fabregat, “End-to-end photoplethysmography (ppg) based biometric authentication by using convolutional neural networks,” in *2018 26th European Signal Processing Conference (EUSIPCO)*, 09 2018, pp. 538–542. 1, 2, 3
- [2] Y. Y. Gu, Y. Zhang, and Y. T. Zhang, “A novel biometric approach in human verification by photoplethysmographic signals,” in *4th International IEEE EMBS Special Topic Conference on Information Technology Applications in Biomedicine, 2003.*, April 2003, pp. 13–14. 1, 2
- [3] Petros Spachos, Jiexin Gao, and Dimitrios Hatzinakos, “Feasibility study of photoplethysmographic signals for biometric identification,” in *2011 17th International Conference on Digital Signal Processing (DSP)*, July 2011, pp. 1–5. 1, 2
- [4] A.V.J. Challoner and C.A. Ramsay, “A photoelectric plethysmograph for the measurement of cutaneous blood flow,” *Physics in Medicine and Biology*, vol. 19, no. 3, pp. 317–328, may 1974. 1
- [5] Zhiling Zhang, Zhouyue Pi, and Benyuan Liu, “Troika: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise,” *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 2, pp. 522–531, Feb 2015. 1
- [6] Yan Michalevsky, Dan Boneh, and Gabi Nakibly, “Gyrophone: Recognizing speech from gyroscope signals,” in *23rd USENIX Security Symposium (USENIX Security 14)*, San Diego, CA, 2014, pp. 1053–1067, USENIX Association. 1, 2
- [7] Aleksandar Matic, Venet Osmani, and Oscar Mayora, “Speech activity detection using accelerometer,” *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 2012, pp. 2112–5, 08 2012. 1, 2
- [8] Seung-Jae Moon and Björn Lindblom, “Two experiments on oxygen consumption during speech production: vocal effort and speaking tempo,” in *Proceedings of the 15th International Congress of Phonetic Sciences*, 2003, pp. 3129–3132. 1, 2
- [9] Jennifer N. John, Conner Galloway, and Alexander Valys, “Deep convolutional neural networks for noise detection in ecgs,” *arXiv preprint arXiv:1810.04122*, 2018. 2, 3
- [10] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014. 2
- [11] Joel Murphy and Yury Gitman, “PulseSensor Open Hardware,” <http://pulsesensor.com/>, 2017, [Online; accessed 19-October-2017]. 2
- [12] Guillem Cortes and Jordi Luque, “PulseID-EUSIPCO,” <https://bitbucket.org/guillemcortes/pulseid-eusipco/>, 2018. 2
- [13] Henk van den Heuvel, Phil Hall, Harald Höge, Asunción Moreno, Antonio Rincon, and Francesco Senia, “SALA II across the finish line: A large collection of mobile telephone speech databases from north and Latin america completed,” in *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal, May 2004, European Language Resources Association (ELRA). 3
- [14] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely, “The Kaldi Speech Recognition Toolkit,” 2011, IEEE Catalog No.: CFP11SRW-USB. 3
- [15] Jordi Luque, Carlos Segura, Ariadna Sánchez, Mart Umbert, and Luis Galindo, “The role of linguistic and prosodic cues on the prediction of self-reported satisfaction in contact centre phone calls,” 08 2017, pp. 2346–2350. 3
- [16] Masaaki Hirotsu “mhiro2”, “Freesound Audio Tagging 2019: Simple 2D-CNN Classifier with PyTorch,” <https://www.kaggle.com/mhiro2/simple-2d-cnn-classifier-with-pytorch/>, 2019. 3
- [17] Ilya Loshchilov and Frank Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” 2016. 3
- [18] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer, “Automatic differentiation in PyTorch,” in *NIPS Autodiff Workshop*, 2017. 3
- [19] José A.R. Fonollosa, “Speech Commands Recognition,” <https://github.com/jarfo/gcommands/>, 2018. 3
- [20] Guillermo Cámara and Jordi Luque, “PulseSpeech,” <https://bitbucket.org/gcambara/pulsespeech/>, 2019. 3