UNIVERSITEIT iYUNIVESITHI STELLENBOSCH UNIVERSITY 100 1918 · 2018

saam vorentoe · masiye phambili · forward together

Deep Neural Network based Cough Detection using Bed-mounted Accelerometer Measurements

Madhurananda Pahar[1], Igor Miranda[1], Andreas Diacon[2], Thomas Niesler[1]

[1]Department of Electrical and Electronic Engineering, University of Stellenbosch, South Africa
[2]TASK Applied Science, Cape Town, South Africa

ICASSP 2021

Overview

- We have performed deep neural network based cough detection using signals captured from the accelerometer in an inexpensive consumer smartphone attached to the patient's bed.

Overview

- We have performed deep neural network based cough detection using signals captured from the accelerometer in an inexpensive consumer smartphone attached to the patient's bed.
- CNN, LSTM and Resnet50 classifiers have been trained and evaluated by leave-one-out cross-validation on a manually-annotated dataset containing the acceleration signals of approximately 6000 cough and 68000 non-cough events from 14 adult male patients in a tuberculosis clinic.

Overview

- We have performed deep neural network based cough detection using signals captured from the accelerometer in an inexpensive consumer smartphone attached to the patient's bed.
- CNN, LSTM and Resnet50 classifiers have been trained and evaluated by leave-one-out cross-validation on a manually-annotated dataset containing the acceleration signals of approximately 6000 cough and 68000 non-cough events from 14 adult male patients in a tuberculosis clinic.
- Although all classifiers successfully detected cough events, the highest AUC over 0.98 has been achieved by the Resnet50.

Motivation

This form of cough monitoring is beneficial because

Motivation

This form of cough monitoring is beneficial because

- It is noninvasive.

This form of cough monitoring is beneficial because

- It is noninvasive.
- It is less intrusive than body-attached accelerometer sensors.

Motivation

This form of cough monitoring is beneficial because

- It is noninvasive.
- It is less intrusive than body-attached accelerometer sensors.
- It sidesteps privacy concerns encountered when using audio for cough detection.

This form of cough monitoring is beneficial because

- It is noninvasive.
- It is less intrusive than body-attached accelerometer sensors.
- It sidesteps privacy concerns encountered when using audio for cough detection.
- Thus it is as an excellent tool for long-term cough monitoring.

- Coughing, distinctive in nature, is the forceful expulsion of air to clear up the airway and a common symptom of more than 100 respiratory disease.

Introduction

- Coughing, distinctive in nature, is the forceful expulsion of air to clear up the airway and a common symptom of more than 100 respiratory disease.
- Previously, body-attached accelerometers have been used to detect movements including coughs with high accuracy.

Introduction

- Coughing, distinctive in nature, is the forceful expulsion of air to clear up the airway and a common symptom of more than 100 respiratory disease.
- Previously, body-attached accelerometers have been used to detect movements including coughs with high accuracy.
- However, attaching an accelerometer to the patient's body is inconvenient and intrusive. Thus, we propose the monitoring of coughing based on the signals obtained from the accelerometer integrated into an inexpensive consumer smartphone firmly attached to the patient's bed, thereby eliminating the need to wear a measuring equipment.

Data Collection

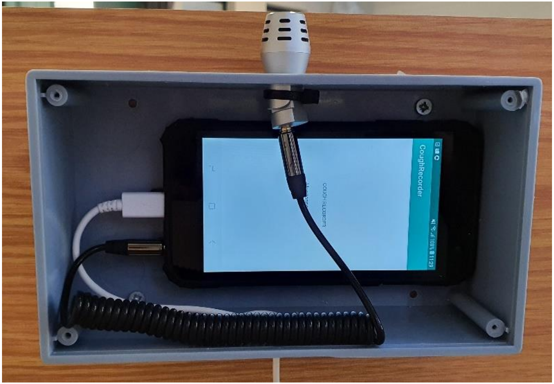

A Samsung Galaxy J4 smartphone with an inbuilt accelerometer connected to an external BOYA BY-MM1 cardioid microphone by a 3.5 mm audio jack

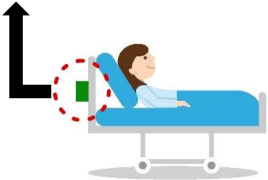

Data Collection

- Data collection was performed at a small 24h TB clinic near Cape Town, South Africa, which can accommodate approximately 10 staff and 30 patients.

Data Collection

- Data collection was performed at a small 24h TB clinic near Cape Town, South Africa, which can accommodate approximately 10 staff and 30 patients.
- Each ward has four beds and can thus accommodate up to four patients at one time.

Data Collection

- Data collection was performed at a small 24h TB clinic near Cape Town, South Africa, which can accommodate approximately 10 staff and 30 patients.
- Each ward has four beds and can thus accommodate up to four patients at one time.
- The 3-axis accelerometer has a sampling frequency of 100 Hz and only magnitudes were recorded.

Data Collection

- Data collection was performed at a small 24h TB clinic near Cape Town, South Africa, which can accommodate approximately 10 staff and 30 patients.
- Each ward has four beds and can thus accommodate up to four patients at one time.
- The 3-axis accelerometer has a sampling frequency of 100 Hz and only magnitudes were recorded.
- In addition, continuous video recordings were made using ceiling-mounted cameras in order to ensure accurate data annotation.

Data Collection

Table: **Ground Truth Dataset Summary:** 'PATIENTS': list of the patients; 'COUGHS': number of cough events; 'NON COUGHS': number of events that are not coughs; 'COUGH TIME': total amount of time (in sec) for cough events; 'NON-COUGH TIME': total amount of time (in sec) for non-cough events.

PATIENTS	COUGHS	NON COUGHS	COUGH TIME	NON-COUGH TIME
Patient 1	88	973	169.16	1660.67
Patient 2	63	1111	117.67	1891.92
Patient 3	469	11025	893.91	18797.32
Patient 4	109	9151	204.06	15596.71
Patient 5	97	7826	188.26	13344.98
Patient 6	192	12437	360.72	21197.35
Patient 7	436	14053	825.23	23953.15
Patient 8	368	2977	702.05	5077.89
Patient 9	2816	3856	5345.27	6569.32
Patient 10	649	2579	1236.84	4400.42
Patient 11	205	527	391.42	901.38
Patient 12	213	323	402.61	547.62
Patient 13	213	712	401.61	1211.75
Patient 14	82	455	158.77	777.64
TOTAL	**6000**	**68005**	**11397.6**	**115928.12**

Data Collection

- Our ground truth dataset contains 6000 coughs and 68000 non-coughs from 14 adult male patients totalling 3.16 and 32.20 hours of data respectively. No other information regarding patients are recorded due to ethical constraints.

Data Collection

- Our ground truth dataset contains 6000 coughs and 68000 non-coughs from 14 adult male patients totalling 3.16 and 32.20 hours of data respectively. No other information regarding patients are recorded due to ethical constraints.
- Coughs are underrepresented in our dataset and to compensate this imbalance, we have applied SMOTE data balancing during training.

Cough Audio & Accelerometer Signal

- Coughs are very distinctive in nature.
- The following figure shows the audio and accelerometer signals two consecutive coughs from a patient and the unique patterns in both audio and accelerometer measurements.

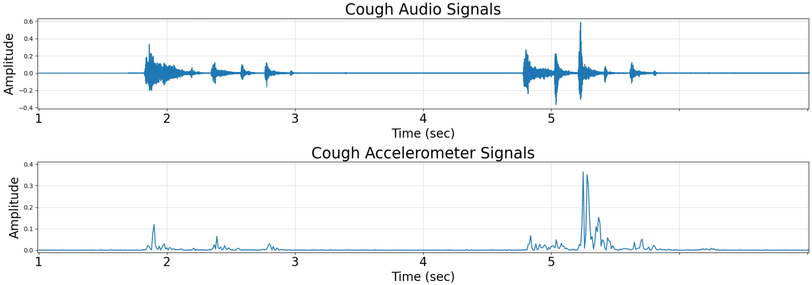

Cough & Non-cough Accelerometer Signal

- Example accelerometer magnitudes for a cough and a non-cough event

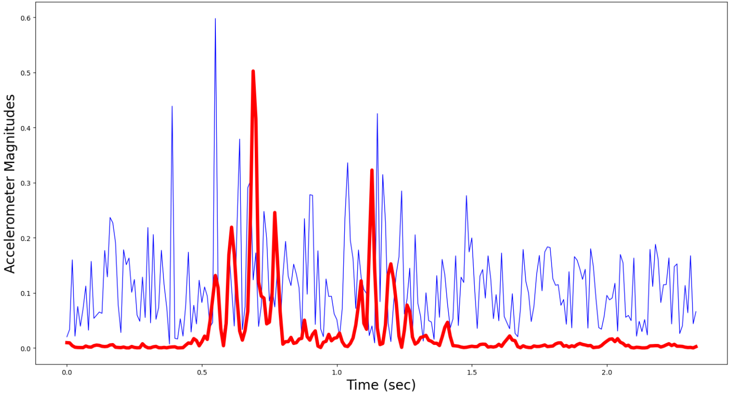

Feature Extraction

- Power spectra, root mean square amplitude, kurtosis, moving average and crest factor are extracted from the acceleration signals as features for cough detection.

Feature Extraction

- Power spectra, root mean square amplitude, kurtosis, moving average and crest factor are extracted from the acceleration signals as features for cough detection.
Hyperparameter Description	**Range**
Size of frames in samples in which cough is segmented (Ψ)	2^k where $k = 4, 5, 6$
Number of segments in which frames were grouped (C)	5, 10
- The input feature matrix has the dimension of $(C, \frac{\Psi}{2} + 5)$ and power spectra have $(\frac{\Psi}{2} + 1)$ coefficients.

Classifier Training

- Only CNN and LSTM has their hyperparameters optimised. The default 50 layer residual architecture used as Resnet50.

Hyperparameters	**Range**
No. of Conv filters (α_1)	3×2^k where $k = 3, 4, 5$
kernel size (α_2)	2 and 3
Dropout rate (α_3)	0.1 to 0.5 in steps of 0.2
Dense layer size (α_4)	2^k where $k = 4, 5$
LSTM units (β_1)	2^k where $k = 6, 7, 8$
Learning rate (β_2)	10^k where $k = -2, -3, -4$
Batch Size (ξ_1)	2^k where $k = 6, 7, 8$
No. of epochs (ξ_2)	10 to 200 in steps of 20

-

Classifier Training

- CNN architecture showing hyperparameters

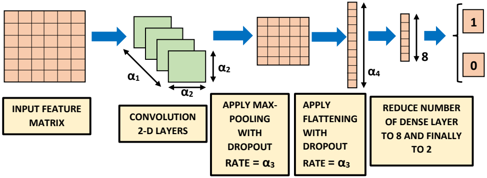

Classifier Training

- CNN architecture showing hyperparameters

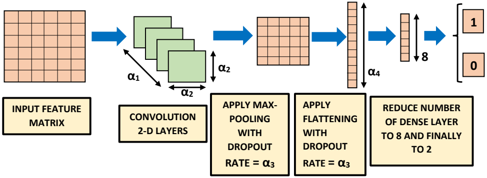

- LSTM architecture showing hyperparameters

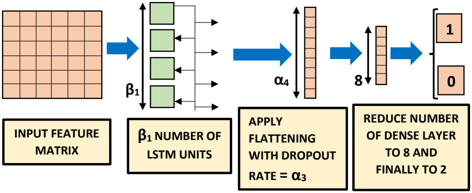

Results

- The best results from these three classifiers:

Clas- -sifier	Frame (Ψ)	Seg (C)	Mean Spec	Mean Sens	Mean Accuracy	Mean AUC
CNN	64	10	91%	80%	86%	0.9499
LSTM	32	10	86%	93%	89%	0.9572
Resnet50	*32*	*10*	*94%*	*99%*	*97%*	*0.9888*

Results

- The best results from these three classifiers:

Clas- -sifier	Frame (Ψ)	Seg (C)	Mean Spec	Mean Sens	Mean Accuracy	Mean AUC
CNN	64	10	91%	80%	86%	0.9499
LSTM	32	10	86%	93%	89%	0.9572
Resnet50	*32*	*10*	*94%*	*99%*	*97%*	*0.9888*

- AUC values are averaged over 14 leave-one-patient-out cross-validation folds during hyperparameter optimisation.

- The best results from these three classifiers:

Clas- -sifier	**Frame (Ψ)**	**Seg (C)**	**Mean Spec**	**Mean Sens**	**Mean Accuracy**	**Mean AUC**
CNN	64	10	91%	80%	86%	0.9499
LSTM	32	10	86%	93%	89%	0.9572
Resnet50	*32*	*10*	*94%*	*99%*	*97%*	*0.9888*

- AUC values are averaged over 14 leave-one-patient-out cross-validation folds during hyperparameter optimisation.
- The best performance is achieved by the Resnet50 architecture, with an AUC of 0.9888 after 50 epochs at training when using 32 sample (320 msec) long frames and 10 segments.

Results

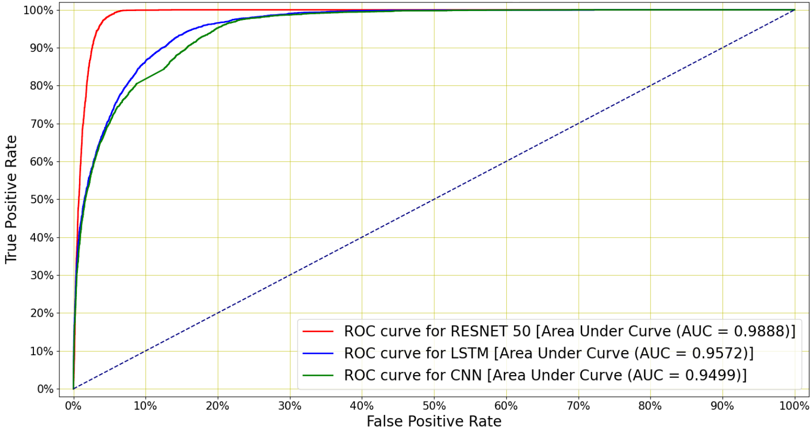

Figure: **Mean ROC curves**, for the best performing classifiers. Resnet50 outperforms the LSTM and CNN over a wide range of operating points.

Conclusion

- A deep neural network based cough detector is able to accurately discriminate between the accelerometer signals due to coughing and due to other movements.

Conclusion

- A deep neural network based cough detector is able to accurately discriminate between the accelerometer signals due to coughing and due to other movements.
- The accelerometer signals were acquired using an integrated accelerometer inside an inexpensive consumer smartphone attached to the patient's bed.

Conclusion

- A deep neural network based cough detector is able to accurately discriminate between the accelerometer signals due to coughing and due to other movements.
- The accelerometer signals were acquired using an integrated accelerometer inside an inexpensive consumer smartphone attached to the patient's bed.
- Although all three classifiers have been successful in detecting cough events, the highest AUC of 0.98 was achieved by the Resnet50.

Conclusion

- A deep neural network based cough detector is able to accurately discriminate between the accelerometer signals due to coughing and due to other movements.
- The accelerometer signals were acquired using an integrated accelerometer inside an inexpensive consumer smartphone attached to the patient's bed.
- Although all three classifiers have been successful in detecting cough events, the highest AUC of 0.98 was achieved by the Resnet50.
- This type of monitoring is less intrusive than approaches in which the accelerometer is attached to the body.

Conclusion

- A deep neural network based cough detector is able to accurately discriminate between the accelerometer signals due to coughing and due to other movements.
- The accelerometer signals were acquired using an integrated accelerometer inside an inexpensive consumer smartphone attached to the patient's bed.
- Although all three classifiers have been successful in detecting cough events, the highest AUC of 0.98 was achieved by the Resnet50.
- This type of monitoring is less intrusive than approaches in which the accelerometer is attached to the body.
- It therefore presents a promising approach to the long-term cough monitoring, which can be of practical use in monitoring the recovery process of patients, for example in the clinic where the data was collected.